```
commit dc073f193b70176b16ae3e6e8afccee07a13df90
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Mon Oct 7 19:01:58 2019 +0200

    Linux 5.3.5

commit d51efd6f9f03cdac21a83bd377de84d9c897a7f9
Author: Eric Biggers <ebiggers@google.com>
Date:   Wed Aug 21 22:16:33 2019 -0700

    vfs: set fs_context::user_ns for reconfigure

    commit 1dd9bc08cf1420d466dd8dcfcc233777e61ca5d2 upstream.

    fs_context::user_ns is used by fuse_parse_param(), even during remount,
    so it needs to be set to the existing value for reconfigure.

    Reproducer:

            #include <fcntl.h>
            #include <sys/mount.h>

            int main()
            {
                    char opts[128];
                    int fd = open("/dev/fuse", O_RDWR);

                    sprintf(opts, "fd=%d,rootmode=040000,user_id=0,group_id=0", fd);
                    mkdir("mnt", 0777);
                    mount("foo",  "mnt", "fuse.foo", 0, opts);
                    mount("foo", "mnt", "fuse.foo", MS_REMOUNT, opts);
            }

    Crash:
            BUG: kernel NULL pointer dereference, address: 0000000000000000
            #PF: supervisor read access in kernel mode
            #PF: error_code(0x0000) - not-present page
            PGD 0 P4D 0
            Oops: 0000 [#1] SMP
            CPU: 0 PID: 129 Comm: syz_make_kuid Not tainted 5.3.0-rc5-next-20190821 #3
            Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.12.0-
20181126_142135-anatol 04/01/2014
            RIP: 0010:map_id_range_down+0xb/0xc0 kernel/user_namespace.c:291
            [...]
            Call Trace:
             map_id_down kernel/user_namespace.c:312 [inline]
             make_kuid+0xe/0x10 kernel/user_namespace.c:389
             fuse_parse_param+0x116/0x210 fs/fuse/inode.c:523
             vfs_parse_fs_param+0xdb/0x1b0 fs/fs_context.c:145
             vfs_parse_fs_string+0x6a/0xa0 fs/fs_context.c:188
             generic_parse_monolithic+0x85/0xc0 fs/fs_context.c:228
             parse_monolithic_mount_data+0x1b/0x20 fs/fs_context.c:708
             do_remount fs/namespace.c:2525 [inline]
             do_mount+0x39a/0xa60 fs/namespace.c:3107
             ksys_mount+0x7d/0xd0 fs/namespace.c:3325
             __do_sys_mount fs/namespace.c:3339 [inline]
             __se_sys_mount fs/namespace.c:3336 [inline]
             __x64_sys_mount+0x20/0x30 fs/namespace.c:3336
             do_syscall_64+0x4a/0x1a0 arch/x86/entry/common.c:290
             entry_SYSCALL_64_after_hwframe+0x49/0xbe

    Reported-by: syzbot+7d6a57304857423318a5@syzkaller.appspotmail.com
    Fixes: 408cbe695350 ("vfs: Convert fuse to use the new mount API")
    Cc: David Howells <dhowells@redhat.com>
```

```
    Cc: Miklos Szeredi <miklos@szeredi.hu>
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Reviewed-by: David Howells <dhowells@redhat.com>
    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9edc7da660df9b699556650d852e4f1f8f7dc739
Author: Bharath Vedartham <linux.bhar@gmail.com>
Date:   Thu May 23 01:15:19 2019 +0530

    9p/cache.c: Fix memory leak in v9fs_cache_session_get_cookie

    commit 962a991c5de18452d6c429d99f3039387cf5cbb0 upstream.

    v9fs_cache_session_get_cookie assigns a random cachetag to v9ses->cachetag,
    if the cachetag is not assigned previously.

    v9fs_random_cachetag allocates memory to v9ses->cachetag with kmalloc and uses
    scnprintf to fill it up with a cachetag.

    But if scnprintf fails, v9ses->cachetag is not freed in the current
    code causing a memory leak.

    Fix this by freeing v9ses->cachetag it v9fs_random_cachetag fails.

    This was reported by syzbot, the link to the report is below:
    https://syzkaller.appspot.com/bug?id=f012bdf297a7a4c860c38a88b44fbee43fd9bbf3

    Link: http://lkml.kernel.org/r/20190522194519.GA5313@bharath12345-Inspiron-5559
    Reported-by: syzbot+3a030a73b6c1e9833815@syzkaller.appspotmail.com
    Signed-off-by: Bharath Vedartham <linux.bhar@gmail.com>
    Signed-off-by: Dominique Martinet <dominique.martinet@cea.fr>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7c8174518d66d8bc95a759db786456aa2abf7fdb
Author: Wanpeng Li <wanpengli@tencent.com>
Date:   Mon Sep 16 15:42:32 2019 +0800

    KVM: hyperv: Fix Direct Synthetic timers assert an interrupt w/o lapic_in_kernel

    commit a073d7e3ad687a7ef32b65affe80faa7ce89bf92 upstream.

    Reported by syzkaller:

            kasan: GPF could be caused by NULL-ptr deref or user memory access
            general protection fault: 0000 [#1] PREEMPT SMP KASAN
            RIP: 0010:__apic_accept_irq+0x46/0x740 arch/x86/kvm/lapic.c:1029
            Call Trace:
            kvm_apic_set_irq+0xb4/0x140 arch/x86/kvm/lapic.c:558
            stimer_notify_direct arch/x86/kvm/hyperv.c:648 [inline]
            stimer_expiration arch/x86/kvm/hyperv.c:659 [inline]
            kvm_hv_process_stimers+0x594/0x1650 arch/x86/kvm/hyperv.c:686
            vcpu_enter_guest+0x2b2a/0x54b0 arch/x86/kvm/x86.c:7896
            vcpu_run+0x393/0xd40 arch/x86/kvm/x86.c:8152
            kvm_arch_vcpu_ioctl_run+0x636/0x900 arch/x86/kvm/x86.c:8360
            kvm_vcpu_ioctl+0x6cf/0xaf0 arch/x86/kvm/../../../virt/kvm/kvm_main.c:2765

    The testcase programs HV_X64_MSR_STIMERn_CONFIG/HV_X64_MSR_STIMERn_COUNT,
    in addition, there is no lapic in the kernel, the counters value are small
    enough in order that kvm_hv_process_stimers() inject this already-expired
    timer interrupt into the guest through lapic in the kernel which triggers
    the NULL deferencing. This patch fixes it by don't advertise direct mode
    synthetic timers and discarding the inject when lapic is not in kernel.

    syzkaller source: https://syzkaller.appspot.com/x/repro.c?x=1752fe0a600000
```

```
    Reported-by: syzbot+dff25ee91f0c7d5c1695@syzkaller.appspotmail.com
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Cc: Radim Krčmář <rkrcmar@redhat.com>
    Signed-off-by: Wanpeng Li <wanpengli@tencent.com>
    Reviewed-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 06d8b3a3a03e7090bc182dbe10bef53d69c46fba
Author: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
Date:   Wed Sep 25 16:47:33 2019 -0700

    kexec: bail out upon SIGKILL when allocating memory.

    commit 7c3a6aedcd6aae0a32a527e68669f7dd667492d1 upstream.

    syzbot found that a thread can stall for minutes inside kexec_load() after
    that thread was killed by SIGKILL [1].  It turned out that the reproducer
    was trying to allocate 2408MB of memory using kimage_alloc_page() from
    kimage_load_normal_segment().  Let's check for SIGKILL before doing memory
    allocation.

    [1] https://syzkaller.appspot.com/bug?id=a0e3436829698d5824231251fad9d8e998f94f5e

    Link: http://lkml.kernel.org/r/993c9185-d324-2640-d061-bed2dd18b1f7@I-love.SAKURA.ne.jp
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Reported-by: syzbot <syzbot+8ab2d0f39fb79fe6ca40@syzkaller.appspotmail.com>
    Cc: Eric Biederman <ebiederm@xmission.com>
    Reviewed-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f2757aaad082ca0a8e3add267130be77725af9b4
Author: Andrey Konovalov <andreyknvl@google.com>
Date:   Mon Jul 29 16:35:01 2019 +0300

    NFC: fix attrs checks in netlink interface

    commit 18917d51472fe3b126a3a8f756c6b18085eb8130 upstream.

    nfc_genl_deactivate_target() relies on the NFC_ATTR_TARGET_INDEX
    attribute being present, but doesn't check whether it is actually
    provided by the user. Same goes for nfc_genl_fw_download() and
    NFC_ATTR_FIRMWARE_NAME.

    This patch adds appropriate checks.

    Found with syzkaller.

    Signed-off-by: Andrey Konovalov <andreyknvl@google.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f020f5c4ccc16c86f4303e735a6caf215f6706f
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Mon Aug 26 02:41:17 2019 -0400

    dm zoned: fix invalid memory access

    commit 0c8e9c2d668278652af028c3cc068c65f66342f4 upstream.

    Commit 75d66ffb48efb30f2dd42f041ba8b39c5b2bd115 ("dm zoned: properly
```

```
    handle backing device failure") triggers a coverity warning:

commit ccf22db4ac539e0379f5983d568a803f9d8d550f
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Sep 11 19:31:33 2019 +0800

    dm raid: fix updating of max_discard_sectors limit

    commit c8156fc77d0796ba2618936dbb3084e769e916c1 upstream.

    Unit of 'chunk_size' is byte, instead of sector, so fix it by setting
    the queue_limits' max_discard_sectors to rs->md.chunk_sectors.  Also,
    rename chunk_size to chunk_size_bytes.

    Without this fix, too big max_discard_sectors is applied on the request
    queue of dm-raid, finally raid code has to split the bio again.

    This re-split done by raid causes the following nested clone_endio:

    1) one big bio 'A' is submitted to dm queue, and served as the original
    bio

    2) one new bio 'B' is cloned from the original bio 'A', and .map()
    is run on this bio of 'B', and B's original bio points to 'A'

    3) raid code sees that 'B' is too big, and split 'B' and re-submit
    the remainded part of 'B' to dm-raid queue via generic_make_request().

    4) now dm will handle 'B' as new original bio, then allocate a new
    clone bio of 'C' and run .map() on 'C'. Meantime C's original bio
    points to 'B'.

    5) suppose now 'C' is completed by raid directly, then the following
    clone_endio() is called recursively:

            clone_endio(C)
                    ->clone_endio(B)                #B is original bio of 'C'
                            ->bio_endio(A)

    'A' can be big enough to make hundreds of nested clone_endio(), then
    stack can be corrupted easily.

    Fixes: 61697a6abd24a ("dm: eliminate 'split_discard_bios' flag from DM target
interface")
    Cc: stable@vger.kernel.org
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ecf169d7697ad7a1e8a1a77726e9b17a53ed7536
Author: Eric Biggers <ebiggers@google.com>
Date:   Wed Aug 21 22:54:41 2019 -0700

    smack: use GFP_NOFS while holding inode_smack::smk_lock

    commit e5bfad3d7acc5702f32aafeb388362994f4d7bd0 upstream.

    inode_smack::smk_lock is taken during smack_d_instantiate(), which is
    called during a filesystem transaction when creating a file on ext4.
    Therefore to avoid a deadlock, all code that takes this lock must use
    GFP_NOFS, to prevent memory reclaim from waiting for the filesystem
    transaction to complete.

    Reported-by: syzbot+0eefc1e06a77d327a056@syzkaller.appspotmail.com
    Cc: stable@vger.kernel.org
```

```
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Casey Schaufler <casey@schaufler-ca.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 93b497173749b711ae40bd43b2224d95c69355fb
Author: Jann Horn <jannh@google.com>
Date:   Thu Jul 4 20:44:44 2019 +0200

    Smack: Don't ignore other bprm->unsafe flags if LSM_UNSAFE_PTRACE is set

    commit 3675f052b43ba51b99b85b073c7070e083f3e6fb upstream.

    There is a logic bug in the current smack_bprm_set_creds():
    If LSM_UNSAFE_PTRACE is set, but the ptrace state is deemed to be
    acceptable (e.g. because the ptracer detached in the meantime), the other
    ->unsafe flags aren't checked. As far as I can tell, this means that
    something like the following could work (but I haven't tested it):

     - task A: create task B with fork()
     - task B: set NO_NEW_PRIVS
     - task B: install a seccomp filter that makes open() return 0 under some
       conditions
     - task B: replace fd 0 with a malicious library
     - task A: attach to task B with PTRACE_ATTACH
     - task B: execve() a file with an SMACK64EXEC extended attribute
     - task A: while task B is still in the middle of execve(), exit (which
       destroys the ptrace relationship)

    Make sure that if any flags other than LSM_UNSAFE_PTRACE are set in
    bprm->unsafe, we reject the execve().

    Cc: stable@vger.kernel.org
    Fixes: 5663884caab1 ("Smack: unify all ptrace accesses in the smack")
    Signed-off-by: Jann Horn <jannh@google.com>
    Signed-off-by: Casey Schaufler <casey@schaufler-ca.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c5475a65b8688f73979dfc524a502740ec78e502
Author: Vladimir Oltean <olteanv@gmail.com>
Date:   Sun Sep 29 02:37:22 2019 +0300

    net: sched: taprio: Avoid division by zero on invalid link speed

    [ Upstream commit 9a9251a3534745d08a92abfeca0ca467b912b5f6 ]

    The check in taprio_set_picos_per_byte is currently not robust enough
    and will trigger this division by zero, due to e.g. PHYLINK not setting
    kset->base.speed when there is no PHY connected:

    [   27.109992] Division by zero in kernel.
    [   27.113842] CPU: 1 PID: 198 Comm: tc Not tainted 5.3.0-rc5-01246-gc4006b8c2637-dirty
#212
    [   27.121974] Hardware name: Freescale LS1021A
    [   27.126234] [<c03132e0>] (unwind_backtrace) from [<c030d8b8>] (show_stack+0x10/0x14)
    [   27.133938] [<c030d8b8>] (show_stack) from [<c10b21b0>] (dump_stack+0xb0/0xc4)
    [   27.141124] [<c10b21b0>] (dump_stack) from [<c10af97c>] (Ldiv0_64+0x8/0x18)
    [   27.148052] [<c10af97c>] (Ldiv0_64) from [<c0700260>] (div64_u64+0xcc/0xf0)
    [   27.154978] [<c0700260>] (div64_u64) from [<c07002d0>] (div64_s64+0x4c/0x68)
    [   27.161993] [<c07002d0>] (div64_s64) from [<c0f3d890>]
(taprio_set_picos_per_byte+0xe8/0xf4)
    [   27.170388] [<c0f3d890>] (taprio_set_picos_per_byte) from [<c0f3f614>]
(taprio_change+0x668/0xcec)
    [   27.179302] [<c0f3f614>] (taprio_change) from [<c0f2bc24>] (qdisc_create+0x1fc/0x4f4)
    [   27.187091] [<c0f2bc24>] (qdisc_create) from [<c0f2c0c8>]
(tc_modify_qdisc+0x1ac/0x6f8)
```

```
      [   27.195055] [<c0f2c0c8>] (tc_modify_qdisc) from [<c0ee9604>]
  (rtnetlink_rcv_msg+0x268/0x2dc)
      [   27.203449] [<c0ee9604>] (rtnetlink_rcv_msg) from [<c0f4fef0>]
  (netlink_rcv_skb+0xe0/0x114)
      [   27.211756] [<c0f4fef0>] (netlink_rcv_skb) from [<c0f4f6cc>]
  (netlink_unicast+0x1b4/0x22c)
      [   27.219977] [<c0f4f6cc>] (netlink_unicast) from [<c0f4fa84>]
  (netlink_sendmsg+0x284/0x340)
      [   27.228198] [<c0f4fa84>] (netlink_sendmsg) from [<c0eae5fc>] (sock_sendmsg+0x14/0x24)
      [   27.235988] [<c0eae5fc>] (sock_sendmsg) from [<c0eaedf8>]
  (___sys_sendmsg+0x214/0x228)
      [   27.243863] [<c0eaedf8>] (___sys_sendmsg) from [<c0eb015c>] (__sys_sendmsg+0x50/0x8c)
      [   27.251652] [<c0eb015c>] (__sys_sendmsg) from [<c0301000>]
  (ret_fast_syscall+0x0/0x54)
      [   27.259524] Exception stack(0xe8045fa8 to 0xe8045ff0)
      [   27.264546] 5fa0:                   b6f608c8 000000f8 00000003 bed7e2f0 00000000
  00000000
      [   27.272681] 5fc0: b6f608c8 000000f8 004ce54c 00000128 5d3ce8c7 00000000 00000026
  00505c9c
      [   27.280812] 5fe0: 00000070 bed7e298 004ddd64 b6dd1e64

      Russell King points out that the ethtool API says zero is a valid return
      value of __ethtool_get_link_ksettings:

         * If it is enabled then they are read-only; if the link
         * is up they represent the negotiated link mode; if the link is down,
         * the speed is 0, %SPEED_UNKNOWN or the highest enabled speed and
         * @duplex is %DUPLEX_UNKNOWN or the best enabled duplex mode.

        So, it seems that taprio is not following the API... I'd suggest either
        fixing taprio, or getting agreement to change the ethtool API.

      The chosen path was to fix taprio.

      Fixes: 7b9eba7ba0c1 ("net/sched: taprio: fix picos_per_byte miscalculation")
      Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
      Acked-by: Vinicius Costa Gomes <vinicius.gomes@intel.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 132dfbb395b992e6634d9a5af45e02465db85ad3
  Author: Vladimir Oltean <olteanv@gmail.com>
  Date:   Sun Sep 29 02:39:48 2019 +0300

      net: sched: cbs: Avoid division by zero when calculating the port rate

      [ Upstream commit 83c8c3cf45163f0c823db37be6ab04dfcf8ac751 ]

      As explained in the "net: sched: taprio: Avoid division by zero on
      invalid link speed" commit, it is legal for the ethtool API to return
      zero as a link speed. So guard against it to ensure we don't perform a
      division by zero in kernel.

      Fixes: e0a7683d30e9 ("net/sched: cbs: fix port_rate miscalculation")
      Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
      Acked-by: Vinicius Costa Gomes <vinicius.gomes@intel.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 56844323565c5f1a5731402defab36f0e8ababc0
  Author: Lorenzo Bianconi <lorenzo@kernel.org>
  Date:   Tue Oct 1 10:33:51 2019 +0200

      net: socionext: netsec: always grab descriptor lock
```

```
    [ Upstream commit 55131dec2b1c7417d216f861ea7a29dc7c4d2d20 ]

    Always acquire tx descriptor spinlock even if a xdp program is not loaded
    on the netsec device since ndo_xdp_xmit can run concurrently with
    netsec_netdev_start_xmit and netsec_clean_tx_dring. This can happen
    loading a xdp program on a different device (e.g virtio-net) and
    xdp_do_redirect_map/xdp_do_redirect_slow can redirect to netsec even if
    we do not have a xdp program on it.

    Fixes: ba2b232108d3 ("net: netsec: add XDP support")
    Tested-by: Ilias Apalodimas <ilias.apalodimas@linaro.org>
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Reviewed-by: Ilias Apalodimas <ilias.apalodimas@linaro.org>
    Acked-by: Toke Høiland-Jørgensen <toke@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d03e652a13f11851b363c697163d823d9669a811
Author: Navid Emamdoost <navid.emamdoost@gmail.com>
Date:   Sun Sep 29 01:43:39 2019 +0300

    net: dsa: sja1105: Prevent leaking memory

    [ Upstream commit 68501df92d116b760777a2cfda314789f926476f ]

    In sja1105_static_config_upload, in two cases memory is leaked: when
    static_config_buf_prepare_for_upload fails and when sja1105_inhibit_tx
    fails. In both cases config_buf should be released.

    Fixes: 8aa9ebccae87 ("net: dsa: Introduce driver for NXP SJA1105 5-port L2 switch")
    Fixes: 1a4c69406cc1 ("net: dsa: sja1105: Prevent PHY jabbering during switch reset")
    Signed-off-by: Navid Emamdoost <navid.emamdoost@gmail.com>
    Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b1aea44ff35345598f95e10c99e3d57f8d4907a4
Author: Vladimir Oltean <olteanv@gmail.com>
Date:   Sun Sep 29 01:08:17 2019 +0300

    net: dsa: sja1105: Ensure PTP time for rxtstamp reconstruction is not in the past

    [ Upstream commit b6f2494d311a19b33b19708543e7ef6dea1de459 ]

    Sometimes the PTP synchronization on the switch 'jumps':

      ptp4l[11241.155]: rms    8 max   16 freq -21732 +/-  11 delay   742 +/-   0
      ptp4l[11243.157]: rms    7 max   17 freq -21731 +/-  10 delay   744 +/-   0
      ptp4l[11245.160]: rms 33592410 max 134217731 freq +192422 +/- 8530253 delay   743 +/-
0
      ptp4l[11247.163]: rms 811631 max 964131 freq +10326 +/- 557785 delay   743 +/-   0
      ptp4l[11249.166]: rms 261936 max 533876 freq -304323 +/- 126371 delay   744 +/-   0
      ptp4l[11251.169]: rms 48700 max 57740 freq -20218 +/- 30532 delay   744 +/-   0
      ptp4l[11253.171]: rms 14570 max 30163 freq  -5568 +/- 7563 delay   742 +/-   0
      ptp4l[11255.174]: rms 2914 max 3440 freq -22001 +/- 1667 delay   744 +/-   1
      ptp4l[11257.177]: rms  811 max 1710 freq -22653 +/- 451 delay   744 +/-   1
      ptp4l[11259.180]: rms  177 max  218 freq -21695 +/-  89 delay   741 +/-   0
      ptp4l[11261.182]: rms   45 max   92 freq -21677 +/-  32 delay   742 +/-   0
      ptp4l[11263.186]: rms   14 max   32 freq -21733 +/-  11 delay   742 +/-   0
      ptp4l[11265.188]: rms    9 max   14 freq -21725 +/-  12 delay   742 +/-   0
      ptp4l[11267.191]: rms    9 max   16 freq -21727 +/-  13 delay   742 +/-   0
      ptp4l[11269.194]: rms    6 max   15 freq -21726 +/-   9 delay   743 +/-   0
      ptp4l[11271.197]: rms    8 max   15 freq -21728 +/-  11 delay   743 +/-   0
      ptp4l[11273.200]: rms    6 max   12 freq -21727 +/-   8 delay   743 +/-   0
      ptp4l[11275.202]: rms    9 max   17 freq -21720 +/-  11 delay   742 +/-   0
```

```
       ptp4l[11277.205]: rms    9 max   18 freq -21725 +/-  12 delay   742 +/-   0
```

    Background: the switch only offers partial RX timestamps (24 bits) and
    it is up to the driver to read the PTP clock to fill those timestamps up
    to 64 bits. But the PTP clock readout needs to happen quickly enough (in
    0.135 seconds, in fact), otherwise the PTP clock will wrap around 24
    bits, condition which cannot be detected.

    Looking at the 'max 134217731' value on output line 3, one can see that
    in hex it is 0x8000003. Because the PTP clock resolution is 8 ns,
    that means 0x1000000 in ticks, which is exactly 2^24. So indeed this is
    a PTP clock wraparound, but the reason might be surprising.

    What is going on is that sja1105_tstamp_reconstruct(priv, now, ts)
    expects a "now" time that is later than the "ts" was snapshotted at.
    This, of course, is obvious: we read the PTP time _after_ the partial RX
    timestamp was received. However, the workqueue is processing frames from
    a skb queue and reuses the same PTP time, read once at the beginning.
    Normally the skb queue only contains one frame and all goes well. But
    when the skb queue contains two frames, the second frame that gets
    dequeued might have been partially timestamped by the RX MAC _after_ we
    had read our PTP time initially.

    The code was originally like that due to concerns that SPI access for
    PTP time readout is a slow process, and we are time-constrained anyway
    (aka: premature optimization). But some timing analysis reveals that the
    time spent until the RX timestamp is completely reconstructed is 1 order
    of magnitude lower than the 0.135 s deadline even under worst-case
    conditions. So we can afford to read the PTP time for each frame in the
    RX timestamping queue, which of course ensures that the full PTP time is
    in the partial timestamp's future.

    Fixes: f3097be21bf1 ("net: dsa: sja1105: Add a state machine for RX timestamping")
    Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fe35c0a5f37d8cc071f740d33fe7852c4989008e
Author: Vladimir Oltean <olteanv@gmail.com>
Date: Tue Oct 1 22:07:01 2019 +0300

    ptp_qoriq: Initialize the registers' spinlock before calling ptp_qoriq_settime

    [Upstream commit db34a4714c013b644eec2de0ec81b1f0373b8b93]

    Because ptp_qoriq_settime is being called prior to spin_lock_init, the
    following stack trace can be seen at driver probe time:

    [2.269117] the code is fine but needs lockdep annotation.
    [2.274569] turning off the locking correctness validator.
    [2.280027] CPU: 1 PID: 1 Comm: swapper/0 Not tainted 5.3.0-rc7-01478-g01eaa67a4797
#263
    [2.288073] Hardware name: Freescale LS1021A
    [2.292337] [<c0313cb4>] (unwind_backtrace) from [<c030e11c>] (show_stack+0x10/0x14)
    [2.300045] [<c030e11c>] (show_stack) from [<c1219440>] (dump_stack+0xcc/0xf8)
    [2.307235] [<c1219440>] (dump_stack) from [<c03b9b44>]
(register_lock_class+0x730/0x73c)
    [2.315372] [<c03b9b44>] (register_lock_class) from [<c03b6190>]
(__lock_acquire+0x78/0x270c)
    [2.323856] [<c03b6190>] (__lock_acquire) from [<c03b90cc>] (lock_acquire+0xe0/0x22c)
    [2.331649] [<c03b90cc>] (lock_acquire) from [<c123c310>]
(_raw_spin_lock_irqsave+0x54/0x68)
    [2.340048] [<c123c310>] (_raw_spin_lock_irqsave) from [<c0e73fe4>]
(ptp_qoriq_settime+0x38/0x80)
    [2.348878] [<c0e73fe4>] (ptp_qoriq_settime) from [<c0e746d4>]

```
 (ptp_qoriq_init+0x1f8/0x484)
     [    2.357189] [<c0e746d4>] (ptp_qoriq_init) from [<c0e74aac>]
 (ptp_qoriq_probe+0xd0/0x184)
     [    2.365243] [<c0e74aac>] (ptp_qoriq_probe) from [<c0b0a07c>]
 (platform_drv_probe+0x48/0x9c)
     [    2.373555] [<c0b0a07c>] (platform_drv_probe) from [<c0b07a14>]
 (really_probe+0x1c4/0x400)
     [    2.381779] [<c0b07a14>] (really_probe) from [<c0b07e28>]
 (driver_probe_device+0x78/0x1b8)
     [    2.390003] [<c0b07e28>] (driver_probe_device) from [<c0b081d0>]
 (device_driver_attach+0x58/0x60)
     [    2.398832] [<c0b081d0>] (device_driver_attach) from [<c0b082d4>]
 (__driver_attach+0xfc/0x160)
     [    2.407402] [<c0b082d4>] (__driver_attach) from [<c0b05a84>]
 (bus_for_each_dev+0x68/0xb4)
     [    2.415539] [<c0b05a84>] (bus_for_each_dev) from [<c0b06b68>]
 (bus_add_driver+0x104/0x20c)
     [    2.423763] [<c0b06b68>] (bus_add_driver) from [<c0b0909c>]
 (driver_register+0x78/0x10c)
     [    2.431815] [<c0b0909c>] (driver_register) from [<c030313c>]
 (do_one_initcall+0x8c/0x3ac)
     [    2.439954] [<c030313c>] (do_one_initcall) from [<c1f013f4>]
 (kernel_init_freeable+0x468/0x548)
     [    2.448610] [<c1f013f4>] (kernel_init_freeable) from [<c12344d8>]
 (kernel_init+0x8/0x10c)
     [    2.456745] [<c12344d8>] (kernel_init) from [<c03010b4>] (ret_from_fork+0x14/0x20)
     [    2.464273] Exception stack(0xea89ffb0 to 0xea89fff8)
     [    2.469297] ffa0:                                     00000000 00000000 00000000
 00000000
     [    2.477432] ffc0: 00000000 00000000 00000000 00000000 00000000 00000000 00000000
 00000000
     [    2.485566] ffe0: 00000000 00000000 00000000 00000000 00000013 00000000

     Fixes: ff54571a747b ("ptp_qoriq: convert to use ptp_qoriq_init/free")
     Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 7fccc8d576004e290395d82c30d42978649c0ac0
 Author: Vladimir Oltean <olteanv@gmail.com>
 Date:   Tue Oct 1 21:58:19 2019 +0300

     net: dsa: sja1105: Fix sleeping while atomic in .port_hwtstamp_set

     [ Upstream commit 3e8db7e56082156a37b71d7334860c10fcea8025 ]

     Currently this stack trace can be seen with CONFIG_DEBUG_ATOMIC_SLEEP=y:

     [   41.568348] BUG: sleeping function called from invalid context at
 kernel/locking/mutex.c:909
     [   41.576757] in_atomic(): 1, irqs_disabled(): 0, pid: 208, name: ptp4l
     [   41.583212] INFO: lockdep is turned off.
     [   41.587123] CPU: 1 PID: 208 Comm: ptp4l Not tainted 5.3.0-rc6-01445-ge950f2d4bc7f-
 dirty #1827
     [   41.599873] [<c0313d7c>] (unwind_backtrace) from [<c030e13c>] (show_stack+0x10/0x14)
     [   41.607584] [<c030e13c>] (show_stack) from [<c1212d50>] (dump_stack+0xd4/0x100)
     [   41.614863] [<c1212d50>] (dump_stack) from [<c037dfc8>] (___might_sleep+0x1c8/0x2b4)
     [   41.622574] [<c037dfc8>] (___might_sleep) from [<c122ea90>] (__mutex_lock+0x48/0xab8)
     [   41.630368] [<c122ea90>] (__mutex_lock) from [<c122f51c>]
 (mutex_lock_nested+0x1c/0x24)
     [   41.638340] [<c122f51c>] (mutex_lock_nested) from [<c0c6fe08>]
 (sja1105_static_config_reload+0x30/0x27c)
     [   41.647779] [<c0c6fe08>] (sja1105_static_config_reload) from [<c0c7015c>]
 (sja1105_hwtstamp_set+0x108/0x1cc)
     [   41.657562] [<c0c7015c>] (sja1105_hwtstamp_set) from [<c0feb650>]
```

```
 (dev_ifsioc+0x18c/0x330)
     [   41.665788] [<c0feb650>] (dev_ifsioc) from [<c0febbd8>] (dev_ioctl+0x320/0x6e8)
     [   41.673064] [<c0febbd8>] (dev_ioctl) from [<c0f8b1f4>] (sock_ioctl+0x334/0x5e8)
     [   41.680340] [<c0f8b1f4>] (sock_ioctl) from [<c05404a8>] (do_vfs_ioctl+0xb0/0xa10)
     [   41.687789] [<c05404a8>] (do_vfs_ioctl) from [<c0540e3c>] (ksys_ioctl+0x34/0x58)
     [   41.695151] [<c0540e3c>] (ksys_ioctl) from [<c0301000>] (ret_fast_syscall+0x0/0x28)
     [   41.702768] Exception stack(0xe8495fa8 to 0xe8495ff0)
     [   41.707796] 5fa0:                   beff4a8c 00000001 00000011 000089b0 beff4a8c
 beff4a80
     [   41.715933] 5fc0: beff4a8c 00000001 0000000c 00000036 b6fa98c8 004e19c1 00000001
 00000000
     [   41.724069] 5fe0: 004dcedc beff4a6c 004c0738 b6e7af4c
     [   41.729860] BUG: scheduling while atomic: ptp4l/208/0x00000002
     [   41.735682] INFO: lockdep is turned off.

     Enabling RX timestamping will logically disturb the fastpath (processing
     of meta frames). Replace bool hwts_rx_en with a bit that is checked
     atomically from the fastpath and temporarily unset from the sleepable
     context during a change of the RX timestamping process (a destructive
     operation anyways, requires switch reset).
     If found unset, the fastpath (net/dsa/tag_sja1105.c) will just drop any
     received meta frame and not take the meta_lock at all.

     Fixes: a602afd200f5 ("net: dsa: sja1105: Expose PTP timestamping ioctls to userspace")
     Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit e74e47c28e20b41311140cd0007c59b3f14e0224
 Author: Dongli Zhang <dongli.zhang@oracle.com>
 Date:   Tue Oct 1 21:56:41 2019 +0800

     xen-netfront: do not use ~0U as error return value for xennet_fill_frags()

     [ Upstream commit a761129e3625688310aecf26e1be9e98e85f8eb5 ]

     xennet_fill_frags() uses ~0U as return value when the sk_buff is not able
     to cache extra fragments. This is incorrect because the return type of
     xennet_fill_frags() is RING_IDX and 0xffffffff is an expected value for
     ring buffer index.

     In the situation when the rsp_cons is approaching 0xffffffff, the return
     value of xennet_fill_frags() may become 0xffffffff which xennet_poll() (the
     caller) would regard as error. As a result, queue->rx.rsp_cons is set
     incorrectly because it is updated only when there is error. If there is no
     error, xennet_poll() would be responsible to update queue->rx.rsp_cons.
     Finally, queue->rx.rsp_cons would point to the rx ring buffer entries whose
     queue->rx_skbs[i] and queue->grant_rx_ref[i] are already cleared to NULL.
     This leads to NULL pointer access in the next iteration to process rx ring
     buffer entries.

     The symptom is similar to the one fixed in
     commit 00b368502d18 ("xen-netfront: do not assume sk_buff_head list is
     empty in error handling").

     This patch changes the return type of xennet_fill_frags() to indicate
     whether it is successful or failed. The queue->rx.rsp_cons will be
     always updated inside this function.

     Fixes: ad4f15dc2c70 ("xen/netfront: don't bug in case of too many frags")
     Signed-off-by: Dongli Zhang <dongli.zhang@oracle.com>
     Reviewed-by: Juergen Gross <jgross@suse.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit fae09a8a15da4e9e8c096e2d7c81dfbcef361ef8
Author: Vladimir Oltean <olteanv@gmail.com>
Date:   Tue Oct 1 21:58:18 2019 +0300

    net: dsa: sja1105: Initialize the meta_lock

    [ Upstream commit d6530e5ad45089c018c3cc5b5957a34721249f6f ]

    Otherwise, with CONFIG_DEBUG_SPINLOCK=y, this stack trace gets printed
    when enabling RX timestamping and receiving a PTP frame:

    [  318.537078] INFO: trying to register non-static key.
    [  318.542040] the code is fine but needs lockdep annotation.
    [  318.547500] turning off the locking correctness validator.
    [  318.552972] CPU: 0 PID: 0 Comm: swapper/0 Not tainted 5.3.0-13257-g0825b0669811-dirty
#1962
    [  318.561283] Hardware name: Freescale LS1021A
    [  318.565566] [<c03144bc>] (unwind_backtrace) from [<c030e164>] (show_stack+0x10/0x14)
    [  318.573289] [<c030e164>] (show_stack) from [<c11b9f50>] (dump_stack+0xd4/0x100)
    [  318.580579] [<c11b9f50>] (dump_stack) from [<c03b9b40>]
(register_lock_class+0x728/0x734)
    [  318.588731] [<c03b9b40>] (register_lock_class) from [<c03b60c4>]
(__lock_acquire+0x78/0x25cc)
    [  318.597227] [<c03b60c4>] (__lock_acquire) from [<c03b8ef8>] (lock_acquire+0xd8/0x234)
    [  318.605033] [<c03b8ef8>] (lock_acquire) from [<c11db934>] (_raw_spin_lock+0x44/0x54)
    [  318.612755] [<c11db934>] (_raw_spin_lock) from [<c1164370>] (sja1105_rcv+0x1f8/0x4e8)
    [  318.620561] [<c1164370>] (sja1105_rcv) from [<c115d7cc>] (dsa_switch_rcv+0x80/0x204)
    [  318.628283] [<c115d7cc>] (dsa_switch_rcv) from [<c0f58c80>]
(__netif_receive_skb_one_core+0x50/0x6c)
    [  318.637386] [<c0f58c80>] (__netif_receive_skb_one_core) from [<c0f58f04>]
(netif_receive_skb_internal+0xac/0x264)
    [  318.647611] [<c0f58f04>] (netif_receive_skb_internal) from [<c0f59e98>]
(napi_gro_receive+0x1d8/0x338)
    [  318.656887] [<c0f59e98>] (napi_gro_receive) from [<c0c298a4>]
(gfar_clean_rx_ring+0x328/0x724)
    [  318.665472] [<c0c298a4>] (gfar_clean_rx_ring) from [<c0c29e60>]
(gfar_poll_rx_sq+0x34/0x94)
    [  318.673795] [<c0c29e60>] (gfar_poll_rx_sq) from [<c0f5b40c>]
(net_rx_action+0x128/0x4f8)
    [  318.681860] [<c0f5b40c>] (net_rx_action) from [<c03022f0>] (__do_softirq+0x148/0x5ac)
    [  318.689666] [<c03022f0>] (__do_softirq) from [<c0355af4>] (irq_exit+0x160/0x170)
    [  318.697040] [<c0355af4>] (irq_exit) from [<c03c6818>] (__handle_domain_irq+0x60/0xb4)
    [  318.704847] [<c03c6818>] (__handle_domain_irq) from [<c07e9440>]
(gic_handle_irq+0x58/0x9c)
    [  318.713172] [<c07e9440>] (gic_handle_irq) from [<c0301a70>] (__irq_svc+0x70/0x98)
    [  318.720622] Exception stack(0xc2001f18 to 0xc2001f60)
    [  318.725656] 1f00:                                                     00000001
00000006
    [  318.733805] 1f20: 00000000 c20165c0 ffffe000 c2010cac c2010cf4 00000001 00000000
c2010c88
    [  318.741955] 1f40: c1f7a5a8 00000000 00000000 c2001f68 c03ba140 c030a288 200e0013
ffffffff
    [  318.750110] [<c0301a70>] (__irq_svc) from [<c030a288>] (arch_cpu_idle+0x24/0x3c)
    [  318.757486] [<c030a288>] (arch_cpu_idle) from [<c038a480>] (do_idle+0x1b8/0x2a4)
    [  318.764859] [<c038a480>] (do_idle) from [<c038a94c>] (cpu_startup_entry+0x18/0x1c)
    [  318.772407] [<c038a94c>] (cpu_startup_entry) from [<c1e00f10>]
(start_kernel+0x4cc/0x4fc)

    Fixes: 844d7edc6a34 ("net: dsa: sja1105: Add a global sja1105_tagger_data structure")
    Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c54182886b93d90279f56fed66862b6975580a6f
Author: Dotan Barak <dotanb@dev.mellanox.co.il>
```

```
Date:   Tue Oct 1 10:21:02 2019 -0700

    net/rds: Fix error handling in rds_ib_add_one()

    [ Upstream commit d64bf89a75b65f83f06be9fb8f978e60d53752db ]

    rds_ibdev:ipaddr_list and rds_ibdev:conn_list are initialized
    after allocation some resources such as protection domain.
    If allocation of such resources fail, then these uninitialized
    variables are accessed in rds_ib_dev_free() in failure path. This
    can potentially crash the system. The code has been updated to
    initialize these variables very early in the function.

    Signed-off-by: Dotan Barak <dotanb@dev.mellanox.co.il>
    Signed-off-by: Sudhakar Dindukurti <sudhakar.dindukurti@oracle.com>
    Acked-by: Santosh Shilimkar <santosh.shilimkar@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ba317ce5f227e0978da96207f99dfd9a8c5a418e
Author: Josh Hunt <johunt@akamai.com>
Date:   Wed Oct 2 13:29:23 2019 -0400

    udp: only do GSO if # of segs > 1

    [ Upstream commit 4094871db1d65810acab3d57f6089aa39ef7f648 ]

    Prior to this change an application sending <= 1MSS worth of data and
    enabling UDP GSO would fail if the system had SW GSO enabled, but the
    same send would succeed if HW GSO offload is enabled. In addition to this
    inconsistency the error in the SW GSO case does not get back to the
    application if sending out of a real device so the user is unaware of this
    failure.

    With this change we only perform GSO if the # of segments is > 1 even
    if the application has enabled segmentation. I've also updated the
    relevant udpgso selftests.

    Fixes: bec1f6f69736 ("udp: generate gso with UDP_SEGMENT")
    Signed-off-by: Josh Hunt <johunt@akamai.com>
    Reviewed-by: Willem de Bruijn <willemb@google.com>
    Reviewed-by: Alexander Duyck <alexander.h.duyck@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 673f813e17b080f86c4199af9ec112ffcaca0fda
Author: Eric Dumazet <edumazet@google.com>
Date:   Mon Sep 30 15:44:44 2019 -0700

    tcp: adjust rto_base in retransmits_timed_out()

    [ Upstream commit 3256a2d6ab1f71f9a1bd2d7f6f18eb8108c48d17 ]

    The cited commit exposed an old retransmits_timed_out() bug
    which assumed it could call tcp_model_timeout() with
    TCP_RTO_MIN as rto_base for all states.

    But flows in SYN_SENT or SYN_RECV state uses a different
    RTO base (1 sec instead of 200 ms, unless BPF choses
    another value)

    This caused a reduction of SYN retransmits from 6 to 4 with
    the default /proc/sys/net/ipv4/tcp_syn_retries value.

    Fixes: a41e8a88b06e ("tcp: better handle TCP_USER_TIMEOUT in SYN_SENT state")
```

```
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Yuchung Cheng <ycheng@google.com>
    Cc: Marek Majkowski <marek@cloudflare.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ff0523263b300ebdbbe43c284b2a3dcd030fd4ac
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Tue Oct 1 16:28:43 2019 +0200

    net: dsa: rtl8366: Check VLAN ID and not ports

    [ Upstream commit e8521e53cca584ddf8ec4584d3c550a6c65f88c4 ]

    There has been some confusion between the port number and
    the VLAN ID in this driver. What we need to check for
    validity is the VLAN ID, nothing else.

    The current confusion came from assigning a few default
    VLANs for default routing and we need to rewrite that
    properly.

    Instead of checking if the port number is a valid VLAN
    ID, check the actual VLAN IDs passed in to the callback
    one by one as expected.

    Fixes: d8652956cf37 ("net: dsa: realtek-smi: Add Realtek SMI driver")
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9e08b782f94be6ec338774f7a7550b166371d365
Author: Dexuan Cui <decui@microsoft.com>
Date:   Mon Sep 30 18:43:50 2019 +0000

    vsock: Fix a lockdep warning in __vsock_release()

    [ Upstream commit 0d9138ffac24cf8b75366ede3a68c951e6dcc575 ]

    Lockdep is unhappy if two locks from the same class are held.

    Fix the below warning for hyperv and virtio sockets (vmci socket code
    doesn't have the issue) by using lock_sock_nested() when __vsock_release()
    is called recursively:

    ============================================
    WARNING: possible recursive locking detected
    5.3.0+ #1 Not tainted
    --------------------------------------------
    server/1795 is trying to acquire lock:
    ffff8880c5158990 (sk_lock-AF_VSOCK){+.+.}, at: hvs_release+0x10/0x120 [hv_sock]

    but task is already holding lock:
    ffff8880c5158150 (sk_lock-AF_VSOCK){+.+.}, at: __vsock_release+0x2e/0xf0 [vsock]

    other info that might help us debug this:
     Possible unsafe locking scenario:

           CPU0
           ----
      lock(sk_lock-AF_VSOCK);
      lock(sk_lock-AF_VSOCK);

     *** DEADLOCK ***
```

```
      May be due to missing lock nesting notation

     2 locks held by server/1795:
      #0: ffff8880c5d05ff8 (&sb->s_type->i_mutex_key#10){+.+.}, at: __sock_release+0x2d/0xa0
      #1: ffff8880c5158150 (sk_lock-AF_VSOCK){+.+.}, at: __vsock_release+0x2e/0xf0 [vsock]

     stack backtrace:
     CPU: 5 PID: 1795 Comm: server Not tainted 5.3.0+ #1
     Call Trace:
      dump_stack+0x67/0x90
      __lock_acquire.cold.67+0xd2/0x20b
      lock_acquire+0xb5/0x1c0
      lock_sock_nested+0x6d/0x90
      hvs_release+0x10/0x120 [hv_sock]
      __vsock_release+0x24/0xf0 [vsock]
      __vsock_release+0xa0/0xf0 [vsock]
      vsock_release+0x12/0x30 [vsock]
      __sock_release+0x37/0xa0
      sock_close+0x14/0x20
      __fput+0xc1/0x250
      task_work_run+0x98/0xc0
      do_exit+0x344/0xc60
      do_group_exit+0x47/0xb0
      get_signal+0x15c/0xc50
      do_signal+0x30/0x720
      exit_to_usermode_loop+0x50/0xa0
      do_syscall_64+0x24e/0x270
      entry_SYSCALL_64_after_hwframe+0x49/0xbe
     RIP: 0033:0x7f4184e85f31

     Tested-by: Stefano Garzarella <sgarzare@redhat.com>
     Signed-off-by: Dexuan Cui <decui@microsoft.com>
     Reviewed-by: Stefano Garzarella <sgarzare@redhat.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 7cc748b229a3796fae59f8909601fbd7e26c2809
  Author: Josh Hunt <johunt@akamai.com>
  Date:   Wed Oct 2 13:29:22 2019 -0400

     udp: fix gso_segs calculations

     [ Upstream commit 44b321e5020d782ad6e8ae8183f09b163be6e6e2 ]

     Commit dfec0ee22c0a ("udp: Record gso_segs when supporting UDP segmentation offload")
     added gso_segs calculation, but incorrectly got sizeof() the pointer and
     not the underlying data type. In addition let's fix the v6 case.

     Fixes: bec1f6f69736 ("udp: generate gso with UDP_SEGMENT")
     Fixes: dfec0ee22c0a ("udp: Record gso_segs when supporting UDP segmentation offload")
     Signed-off-by: Josh Hunt <johunt@akamai.com>
     Reviewed-by: Alexander Duyck <alexander.h.duyck@linux.intel.com>
     Acked-by: Willem de Bruijn <willemb@google.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit f19de0253a0d5d3e64ec46d4d5e50efc39fa9663
  Author: Tuong Lien <tuong.t.lien@dektech.com.au>
  Date:   Wed Oct 2 18:49:43 2019 +0700

     tipc: fix unlimited bundling of small messages

     [ Upstream commit e95584a889e1902fdf1ded9712e2c3c3083baf96 ]

     We have identified a problem with the "oversubscription" policy in the
```

```
    link transmission code.

    When small messages are transmitted, and the sending link has reached
    the transmit window limit, those messages will be bundled and put into
    the link backlog queue. However, bundles of data messages are counted
    at the 'CRITICAL' level, so that the counter for that level, instead of
    the counter for the real, bundled message's level is the one being
    increased.
    Subsequent, to-be-bundled data messages at non-CRITICAL levels continue
    to be tested against the unchanged counter for their own level, while
    contributing to an unrestrained increase at the CRITICAL backlog level.

    This leaves a gap in congestion control algorithm for small messages
    that can result in starvation for other users or a "real" CRITICAL
    user. Even that eventually can lead to buffer exhaustion & link reset.

    We fix this by keeping a 'target_bskb' buffer pointer at each levels,
    then when bundling, we only bundle messages at the same importance
    level only. This way, we know exactly how many slots a certain level
    have occupied in the queue, so can manage level congestion accurately.

    By bundling messages at the same level, we even have more benefits. Let
    consider this:
    - One socket sends 64-byte messages at the 'CRITICAL' level;
    - Another sends 4096-byte messages at the 'LOW' level;

    When a 64-byte message comes and is bundled the first time, we put the
    overhead of message bundle to it (+ 40-byte header, data copy, etc.)
    for later use, but the next message can be a 4096-byte one that cannot
    be bundled to the previous one. This means the last bundle carries only
    one payload message which is totally inefficient, as for the receiver
    also! Later on, another 64-byte message comes, now we make a new bundle
    and the same story repeats...

    With the new bundling algorithm, this will not happen, the 64-byte
    messages will be bundled together even when the 4096-byte message(s)
    comes in between. However, if the 4096-byte messages are sent at the
    same level i.e. 'CRITICAL', the bundling algorithm will again cause the
    same overhead.

    Also, the same will happen even with only one socket sending small
    messages at a rate close to the link transmit's one, so that, when one
    message is bundled, it's transmitted shortly. Then, another message
    comes, a new bundle is created and so on...

    We will solve this issue radically by another patch.

    Fixes: 365ad353c256 ("tipc: reduce risk of user starvation during link congestion")
    Reported-by: Hoang Le <hoang.h.le@dektech.com.au>
    Acked-by: Jon Maloy <jon.maloy@ericsson.com>
    Signed-off-by: Tuong Lien <tuong.t.lien@dektech.com.au>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 151f7893c364a86c61c86845abbb0ee43987c238
Author: Eric Dumazet <edumazet@google.com>
Date:   Fri Oct 4 10:34:45 2019 -0700

    sch_dsmark: fix potential NULL deref in dsmark_init()

    [ Upstream commit 474f0813a3002cb299bb73a5a93aa1f537a80ca8 ]

    Make sure TCA_DSMARK_INDICES was provided by the user.

    syzbot reported :
```

```
    kasan: CONFIG_KASAN_INLINE enabled
    kasan: GPF could be caused by NULL-ptr deref or user memory access
    general protection fault: 0000 [#1] PREEMPT SMP KASAN
    CPU: 1 PID: 8799 Comm: syz-executor235 Not tainted 5.3.0+ #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
01/01/2011
    RIP: 0010:nla_get_u16 include/net/netlink.h:1501 [inline]
    RIP: 0010:dsmark_init net/sched/sch_dsmark.c:364 [inline]
    RIP: 0010:dsmark_init+0x193/0x640 net/sched/sch_dsmark.c:339
    Code: 85 db 58 0f 88 7d 03 00 00 e8 e9 1a ac fb 48 8b 9d 70 ff ff ff 48 b8 00 00 00 00
00 fc ff df 48 8d 7b 04 48 89 fa 48 c1 ea 03 <0f> b6 14 02 48 89 f8 83 e0 07 83 c0 01 38 d0
7c 08 84 d2 0f 85 ca
    RSP: 0018:ffff88809426f3b8 EFLAGS: 00010247
    RAX: dfffc0000000000 RBX: 0000000000000000 RCX: ffffffff85c6eb09
    RDX: 0000000000000000 RSI: ffffffff85c6eb17 RDI: 0000000000000004
    RBP: ffff88809426f4b0 R08: ffff88808c4085c0 R09: ffffed1015d26159
    R10: ffffed1015d26158 R11: ffff8880ae930ac7 R12: ffff8880a7e96940
    R13: dfffc0000000000 R14: ffff88809426f8c0 R15: 0000000000000000
    FS:  0000000001292880(0000) GS:ffff8880ae900000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 0000000020000080 CR3: 000000008ca1b000 CR4: 00000000001406e0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     qdisc_create+0x4ee/0x1210 net/sched/sch_api.c:1237
     tc_modify_qdisc+0x524/0x1c50 net/sched/sch_api.c:1653
     rtnetlink_rcv_msg+0x463/0xb00 net/core/rtnetlink.c:5223
     netlink_rcv_skb+0x177/0x450 net/netlink/af_netlink.c:2477
     rtnetlink_rcv+0x1d/0x30 net/core/rtnetlink.c:5241
     netlink_unicast_kernel net/netlink/af_netlink.c:1302 [inline]
     netlink_unicast+0x531/0x710 net/netlink/af_netlink.c:1328
     netlink_sendmsg+0x8a5/0xd60 net/netlink/af_netlink.c:1917
     sock_sendmsg_nosec net/socket.c:637 [inline]
     sock_sendmsg+0xd7/0x130 net/socket.c:657
     ___sys_sendmsg+0x803/0x920 net/socket.c:2311
     __sys_sendmsg+0x105/0x1d0 net/socket.c:2356
     __do_sys_sendmsg net/socket.c:2365 [inline]
     __se_sys_sendmsg net/socket.c:2363 [inline]
     __x64_sys_sendmsg+0x78/0xb0 net/socket.c:2363
     do_syscall_64+0xfa/0x760 arch/x86/entry/common.c:290
     entry_SYSCALL_64_after_hwframe+0x49/0xbe
    RIP: 0033:0x440369

    Fixes: 758cc43c6d73 ("[PKT_SCHED]: Fix dsmark to apply changes consistent")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 68aa587290eb3865c629c3e57d2fc604c0be7fd1
Author: Eric Dumazet <edumazet@google.com>
Date:   Thu Sep 26 18:24:43 2019 -0700

    sch_cbq: validate TCA_CBQ_WRROPT to avoid crash

    [ Upstream commit e9789c7cc182484fc031fd88097eb14cb26c4596 ]

    syzbot reported a crash in cbq_normalize_quanta() caused
    by an out of range cl->priority.

    iproute2 enforces this check, but malicious users do not.

    kasan: CONFIG_KASAN_INLINE enabled
    kasan: GPF could be caused by NULL-ptr deref or user memory access
```

```
    general protection fault: 0000 [#1] SMP KASAN PTI
    Modules linked in:
    CPU: 1 PID: 26447 Comm: syz-executor.1 Not tainted 5.3+ #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
01/01/2011
    RIP: 0010:cbq_normalize_quanta.part.0+0x1fd/0x430 net/sched/sch_cbq.c:902
    RSP: 0018:ffff8801a5c333b0 EFLAGS: 00010206
    RAX: 0000000020000003 RBX: 00000000fffffff8 RCX: ffffc9000712f000
    RDX: 00000000000043bf RSI: ffffffff83be8962 RDI: 0000000100000018
    RBP: ffff8801a5c33420 R08: 000000000000003a R09: 0000000000000000
    R10: 0000000000000000 R11: 0000000000000000 R12: 00000000000002ef
    R13: ffff88018da95188 R14: dfffffc0000000000 R15: 0000000000000015
    FS:  00007f37d26b1700(0000) GS:ffff8801dad00000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00000000004c7cec CR3: 00000001bcd0a006 CR4: 00000000001626f0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     [<ffffffff83be9d57>] cbq_normalize_quanta include/net/pkt_sched.h:27 [inline]
     [<ffffffff83be9d57>] cbq_addprio net/sched/sch_cbq.c:1097 [inline]
     [<ffffffff83be9d57>] cbq_set_wrr+0x2d7/0x450 net/sched/sch_cbq.c:1115
     [<ffffffff83bee8a7>] cbq_change_class+0x987/0x225b net/sched/sch_cbq.c:1537
     [<ffffffff83b96985>] tc_ctl_tclass+0x555/0xcd0 net/sched/sch_api.c:2329
     [<ffffffff83a84655>] rtnetlink_rcv_msg+0x485/0xc10 net/core/rtnetlink.c:5248
     [<ffffffff83cadf0a>] netlink_rcv_skb+0x17a/0x460 net/netlink/af_netlink.c:2510
     [<ffffffff83a7db6d>] rtnetlink_rcv+0x1d/0x30 net/core/rtnetlink.c:5266
     [<ffffffff83cac2c6>] netlink_unicast_kernel net/netlink/af_netlink.c:1324 [inline]
     [<ffffffff83cac2c6>] netlink_unicast+0x536/0x720 net/netlink/af_netlink.c:1350
     [<ffffffff83cacd4a>] netlink_sendmsg+0x89a/0xd50 net/netlink/af_netlink.c:1939
     [<ffffffff8399d46e>] sock_sendmsg_nosec net/socket.c:673 [inline]
     [<ffffffff8399d46e>] sock_sendmsg+0x12e/0x170 net/socket.c:684
     [<ffffffff8399f1fd>] ___sys_sendmsg+0x81d/0x960 net/socket.c:2359
     [<ffffffff839a2d05>] __sys_sendmsg+0x105/0x1d0 net/socket.c:2397
     [<ffffffff839a2df9>] SYSC_sendmsg net/socket.c:2406 [inline]
     [<ffffffff839a2df9>] SyS_sendmsg+0x29/0x30 net/socket.c:2404
     [<ffffffff8101ccc8>] do_syscall_64+0x528/0x770 arch/x86/entry/common.c:305
     [<ffffffff84400091>] entry_SYSCALL_64_after_hwframe+0x42/0xb7

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6144d252ffc4cbea998808e206ebce1ffc42e09a
Author: David Howells <dhowells@redhat.com>
Date:   Thu Oct 3 17:44:44 2019 +0100

    rxrpc: Fix rxrpc_recvmsg tracepoint

    [ Upstream commit db9b2e0af605e7c994784527abfd9276cabd718a ]

    Fix the rxrpc_recvmsg tracepoint to handle being called with a NULL call
    parameter.

    Fixes: a25e21f0bcd2 ("rxrpc, afs: Use debug_ids rather than pointers in traces")
    Signed-off-by: David Howells <dhowells@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e0869dc8f1630f642ede13246b99e9fe56a469eb
Author: Reinhard Speyerer <rspmn@arcor.de>
Date:   Thu Oct 3 18:34:39 2019 +0200

    qmi_wwan: add support for Cinterion CLS8 devices
```

```
    [ Upstream commit cf74ac6db25d4002089e85cc623ad149ecc25614 ]

    Add support for Cinterion CLS8 devices.
    Use QMI_QUIRK_SET_DTR as required for Qualcomm MDM9x07 chipsets.

    T:  Bus=01 Lev=03 Prnt=05 Port=01 Cnt=02 Dev#= 25 Spd=480  MxCh= 0
    D:  Ver= 2.00 Cls=00(>ifc ) Sub=00 Prot=00 MxPS=64 #Cfgs=  1
    P:  Vendor=1e2d ProdID=00b0 Rev= 3.18
    S:  Manufacturer=GEMALTO
    S:  Product=USB Modem
    C:* #Ifs= 5 Cfg#= 1 Atr=80 MxPwr=500mA
    I:* If#= 0 Alt= 0 #EPs= 2 Cls=ff(vend.) Sub=42 Prot=01 Driver=(none)
    E:  Ad=01(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    E:  Ad=81(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    I:* If#= 1 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=00 Prot=00 Driver=option
    E:  Ad=83(I) Atr=03(Int.) MxPS=  10 Ivl=32ms
    E:  Ad=82(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    E:  Ad=02(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    I:* If#= 2 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=00 Prot=00 Driver=option
    E:  Ad=85(I) Atr=03(Int.) MxPS=  10 Ivl=32ms
    E:  Ad=84(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    E:  Ad=03(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    I:* If#= 3 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=00 Prot=00 Driver=option
    E:  Ad=87(I) Atr=03(Int.) MxPS=  10 Ivl=32ms
    E:  Ad=86(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    E:  Ad=04(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    I:* If#= 4 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=ff Prot=ff Driver=qmi_wwan
    E:  Ad=89(I) Atr=03(Int.) MxPS=   8 Ivl=32ms
    E:  Ad=88(I) Atr=02(Bulk) MxPS= 512 Ivl=0ms
    E:  Ad=05(O) Atr=02(Bulk) MxPS= 512 Ivl=0ms

    Signed-off-by: Reinhard Speyerer <rspmn@arcor.de>
    Acked-by: Bjørn Mork <bjorn@mork.no>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3c92e26957e3195365a9ec6f056a95aefa38655c
Author: Eric Dumazet <edumazet@google.com>
Date:   Fri Oct 4 11:08:34 2019 -0700

    nfc: fix memory leak in llcp_sock_bind()

    [ Upstream commit a0c2dc1fe63e2869b74c1c7f6a81d1745c8a695d ]

    sysbot reported a memory leak after a bind() has failed.

    While we are at it, abort the operation if kmemdup() has failed.

    BUG: memory leak
    unreferenced object 0xffff888105d83ec0 (size 32):
      comm "syz-executor067", pid 7207, jiffies 4294956228 (age 19.430s)
      hex dump (first 32 bytes):
        00 69 6c 65 20 72 65 61 64 00 6e 65 74 3a 5b 34  .ile read.net:[4
        30 32 36 35 33 33 30 39 37 5d 00 00 00 00 00 00  026533097]......
      backtrace:
        [<0000000036bac473>] kmemleak_alloc_recursive /./include/linux/kmemleak.h:43
[inline]
        [<0000000036bac473>] slab_post_alloc_hook /mm/slab.h:522 [inline]
        [<0000000036bac473>] slab_alloc /mm/slab.c:3319 [inline]
        [<0000000036bac473>] __do_kmalloc /mm/slab.c:3653 [inline]
        [<0000000036bac473>] __kmalloc_track_caller+0x169/0x2d0 /mm/slab.c:3670
        [<000000000cd39d07>] kmemdup+0x27/0x60 /mm/util.c:120
        [<000000008e57e5fc>] kmemdup /./include/linux/string.h:432 [inline]
        [<000000008e57e5fc>] llcp_sock_bind+0x1b3/0x230 /net/nfc/llcp_sock.c:107
```

```
         [<000000009cb0b5d3>] __sys_bind+0x11c/0x140 /net/socket.c:1647
         [<00000000492c3bbc>] __do_sys_bind /net/socket.c:1658 [inline]
         [<00000000492c3bbc>] __se_sys_bind /net/socket.c:1656 [inline]
         [<00000000492c3bbc>] __x64_sys_bind+0x1e/0x30 /net/socket.c:1656
         [<0000000008704b2a>] do_syscall_64+0x76/0x1a0 /arch/x86/entry/common.c:296
         [<000000009f4c57a4>] entry_SYSCALL_64_after_hwframe+0x44/0xa9

    Fixes: 30cc4587659e ("NFC: Move LLCP code to the NFC top level diirectory")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7e2777fd4816cdf6bff5de9e5221514f36dddfbf
Author: Martin KaFai Lau <kafai@fb.com>
Date:   Fri Sep 27 16:00:31 2019 -0700

    net: Unpublish sk from sk_reuseport_cb before call_rcu

    [ Upstream commit 8c7138b33e5c690c308b2a7085f6313fdcb3f616 ]

    The "reuse->sock[]" array is shared by multiple sockets.  The going away
    sk must unpublish itself from "reuse->sock[]" before making call_rcu()
    call.  However, this unpublish-action is currently done after a grace
    period and it may cause use-after-free.

    The fix is to move reuseport_detach_sock() to sk_destruct().
    Due to the above reason, any socket with sk_reuseport_cb has
    to go through the rcu grace period before freeing it.

    It is a rather old bug (~3 yrs).  The Fixes tag is not necessary
    the right commit but it is the one that introduced the SOCK_RCU_FREE
    logic and this fix is depending on it.

    Fixes: a4298e4522d6 ("net: add SOCK_RCU_FREE socket flag")
    Cc: Eric Dumazet <eric.dumazet@gmail.com>
    Suggested-by: Eric Dumazet <eric.dumazet@gmail.com>
    Signed-off-by: Martin KaFai Lau <kafai@fb.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0711307b9b65d13b38fb30e601d5a9f562066ccb
Author: Vladimir Oltean <olteanv@gmail.com>
Date:   Sun Sep 29 02:01:39 2019 +0300

    net: sched: taprio: Fix potential integer overflow in taprio_set_picos_per_byte

    [ Upstream commit 68ce6688a5baefde30914fc07fc27292dbbe8320 ]

    The speed divisor is used in a context expecting an s64, but it is
    evaluated using 32-bit arithmetic.

    To avoid that happening, instead of multiplying by 1,000,000 in the
    first place, simplify the fraction and do a standard 32 bit division
    instead.

    Fixes: f04b514c0ce2 ("taprio: Set default link speed to 10 Mbps in
taprio_set_picos_per_byte")
    Reported-by: Gustavo A. R. Silva <gustavo@embeddedor.com>
    Suggested-by: Eric Dumazet <eric.dumazet@gmail.com>
    Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
    Acked-by: Vinicius Costa Gomes <vinicius.gomes@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 8dde3868b536a684d7ff5c2f9da9dd48bd71baa0
Author: Navid Emamdoost <navid.emamdoost@gmail.com>
Date:   Fri Oct 4 15:24:39 2019 -0500

    net: qlogic: Fix memory leak in ql_alloc_large_buffers

    [ Upstream commit 1acb8f2a7a9f10543868ddd737e37424d5c36cf4 ]

    In ql_alloc_large_buffers, a new skb is allocated via netdev_alloc_skb.
    This skb should be released if pci_dma_mapping_error fails.

    Fixes: 0f8ab89e825f ("qla3xxx: Check return code from pci_map_single() in
ql_release_to_lrg_buf_free_list(), ql_populate_free_queue(), ql_alloc_large_buffers(), and
ql3xxx_send()")
    Signed-off-by: Navid Emamdoost <navid.emamdoost@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2b162077ebe6f34257c130f9c8864facdab65ae
Author: Paolo Abeni <pabeni@redhat.com>
Date:   Fri Oct 4 15:11:17 2019 +0200

    net: ipv4: avoid mixed n_redirects and rate_tokens usage

    [ Upstream commit b406472b5ad79ede8d10077f0c8f05505ace8b6d ]

    Since commit c09551c6ff7f ("net: ipv4: use a dedicated counter
    for icmp_v4 redirect packets") we use 'n_redirects' to account
    for redirect packets, but we still use 'rate_tokens' to compute
    the redirect packets exponential backoff.

    If the device sent to the relevant peer any ICMP error packet
    after sending a redirect, it will also update 'rate_token' according
    to the leaking bucket schema; typically 'rate_token' will raise
    above BITS_PER_LONG and the redirect packets backoff algorithm
    will produce undefined behavior.

    Fix the issue using 'n_redirects' to compute the exponential backoff
    in ip_rt_send_redirect().

    Note that we still clear rate_tokens after a redirect silence period,
    to avoid changing an established behaviour.

    The root cause predates git history; before the mentioned commit in
    the critical scenario, the kernel stopped sending redirects, after
    the mentioned commit the behavior more randomic.

    Reported-by: Xiumei Mu <xmu@redhat.com>
    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Fixes: c09551c6ff7f ("net: ipv4: use a dedicated counter for icmp_v4 redirect packets")
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Acked-by: Lorenzo Bianconi <lorenzo.bianconi@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7924d9fec76d023014ce2aab567aaf12e19737e4
Author: David Ahern <dsahern@gmail.com>
Date:   Fri Oct 4 08:03:09 2019 -0700

    ipv6: Handle missing host route in __ipv6_ifa_notify

    [ Upstream commit 2d819d250a1393a3e725715425ab70a0e0772a71 ]

    Rajendra reported a kernel panic when a link was taken down:
```

```
        [ 6870.263084] BUG: unable to handle kernel NULL pointer dereference at
 00000000000000a8
        [ 6870.271856] IP: [<ffffffff8efc5764>] __ipv6_ifa_notify+0x154/0x290

        <snip>

        [ 6870.570501] Call Trace:
        [ 6870.573238] [<ffffffff8efc58c6>] ? ipv6_ifa_notify+0x26/0x40
        [ 6870.579665] [<ffffffff8efc98ec>] ? addrconf_dad_completed+0x4c/0x2c0
        [ 6870.586869] [<ffffffff8efe70c6>] ? ipv6_dev_mc_inc+0x196/0x260
        [ 6870.593491] [<ffffffff8efc9c6a>] ? addrconf_dad_work+0x10a/0x430
        [ 6870.600305] [<ffffffff8f01ade4>] ? __switch_to_asm+0x34/0x70
        [ 6870.606732] [<ffffffff8ea93a7a>] ? process_one_work+0x18a/0x430
        [ 6870.613449] [<ffffffff8ea93d6d>] ? worker_thread+0x4d/0x490
        [ 6870.619778] [<ffffffff8ea93d20>] ? process_one_work+0x430/0x430
        [ 6870.626495] [<ffffffff8ea99dd9>] ? kthread+0xd9/0xf0
        [ 6870.632145] [<ffffffff8f01ade4>] ? __switch_to_asm+0x34/0x70
        [ 6870.638573] [<ffffffff8ea99d00>] ? kthread_park+0x60/0x60
        [ 6870.644707] [<ffffffff8f01ae77>] ? ret_from_fork+0x57/0x70
        [ 6870.650936] Code: 31 c0 31 d2 41 b9 20 00 08 02 b9 09 00 00 0
```

    addrconf_dad_work is kicked to be scheduled when a device is brought
    up. There is a race between addrcond_dad_work getting scheduled and
    taking the rtnl lock and a process taking the link down (under rtnl).
    The latter removes the host route from the inet6_addr as part of
    addrconf_ifdown which is run for NETDEV_DOWN. The former attempts
    to use the host route in __ipv6_ifa_notify. If the down event removes
    the host route due to the race to the rtnl, then the BUG listed above
    occurs.

    Since the DAD sequence can not be aborted, add a check for the missing
    host route in __ipv6_ifa_notify. The only way this should happen is due
    to the previously mentioned race. The host route is created when the
    address is added to an interface; it is only removed on a down event
    where the address is kept. Add a warning if the host route is missing
    AND the device is up; this is a situation that should never happen.

    Fixes: f1705ec197e7 ("net: ipv6: Make address flushing on ifdown optional")
    Reported-by: Rajendra Dendukuri <rajendra.dendukuri@broadcom.com>
    Signed-off-by: David Ahern <dsahern@gmail.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 81d14702ad7cc7dc6e6e65d5c18160882ca351df
Author: Eric Dumazet <edumazet@google.com>
Date: Wed Oct 2 09:38:55 2019 -0700

    ipv6: drop incoming packets having a v4mapped source address

    [Upstream commit 6af1799aaf3f1bc8defedddfa00df3192445bbf3]

    This began with a syzbot report. syzkaller was injecting
    IPv6 TCP SYN packets having a v4mapped source address.

    After an unsuccessful 4-tuple lookup, TCP creates a request
    socket (SYN_RECV) and calls reqsk_queue_hash_req()

    reqsk_queue_hash_req() calls sk_ehashfn(sk)

    At this point we have AF_INET6 sockets, and the heuristic
    used by sk_ehashfn() to either hash the IPv4 or IPv6 addresses
    is to use ipv6_addr_v4mapped(&sk->sk_v6_daddr)

    For the particular spoofed packet, we end up hashing V4 addresses

```
    which were not initialized by the TCP IPv6 stack, so KMSAN fired
    a warning.

    I first fixed sk_ehashfn() to test both source and destination addresses,
    but then faced various problems, including user-space programs
    like packetdrill that had similar assumptions.

    Instead of trying to fix the whole ecosystem, it is better
    to admit that we have a dual stack behavior, and that we
    can not build linux kernels without V4 stack anyway.

    The dual stack API automatically forces the traffic to be IPv4
    if v4mapped addresses are used at bind() or connect(), so it makes
    no sense to allow IPv6 traffic to use the same v4mapped class.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Florian Westphal <fw@strlen.de>
    Cc: Hannes Frederic Sowa <hannes@stressinduktion.org>
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b67e416a6495db3f2cad884bc0433b6d43e838e3
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Sep 30 17:12:41 2019 +0200

    hso: fix NULL-deref on tty open

    [ Upstream commit 8353da9fa69722b54cba82b2ec740afd3d438748 ]

    Fix NULL-pointer dereference on tty open due to a failure to handle a
    missing interrupt-in endpoint when probing modem ports:

            BUG: kernel NULL pointer dereference, address: 0000000000000006
            ...
            RIP: 0010:tiocmget_submit_urb+0x1c/0xe0 [hso]
            ...
            Call Trace:
            hso_start_serial_device+0xdc/0x140 [hso]
            hso_serial_open+0x118/0x1b0 [hso]
            tty_open+0xf1/0x490

    Fixes: 542f54823614 ("tty: Modem functions for the HSO driver")
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 76802e205fee12622a3a7ccf47c872b83c92a052
Author: Haishuang Yan <yanhaishuang@cmss.chinamobile.com>
Date:   Fri Sep 27 14:58:20 2019 +0800

    erspan: remove the incorrect mtu limit for erspan

    [ Upstream commit 0e141f757b2c78c983df893e9993313e2dc21e38 ]

    erspan driver calls ether_setup(), after commit 61e84623ace3
    ("net: centralize net_device min/max MTU checking"), the range
    of mtu is [min_mtu, max_mtu], which is [68, 1500] by default.

    It causes the dev mtu of the erspan device to not be greater
    than 1500, this limit value is not correct for ipgre tap device.

    Tested:
    Before patch:
```

```
    # ip link set erspan0 mtu 1600
    Error: mtu greater than device maximum.
    After patch:
    # ip link set erspan0 mtu 1600
    # ip -d link show erspan0
    21: erspan0@NONE: <BROADCAST,MULTICAST> mtu 1600 qdisc noop state DOWN
    mode DEFAULT group default qlen 1000
        link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff promiscuity 0 minmtu 68 maxmtu 0

    Fixes: 61e84623ace3 ("net: centralize net_device min/max MTU checking")
    Signed-off-by: Haishuang Yan <yanhaishuang@cmss.chinamobile.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 38baaf9f4cb4e660f7c3e78154f23a74e7d11468
Author: Vishal Kulkarni <vishal@chelsio.com>
Date:   Fri Oct 4 04:06:15 2019 +0530

    cxgb4:Fix out-of-bounds MSI-X info array access

    [ Upstream commit 6b517374f4ea5a3c6e307e1219ec5f35d42e6d00 ]

    When fetching free MSI-X vectors for ULDs, check for the error code
    before accessing MSI-X info array. Otherwise, an out-of-bounds access is
    attempted, which results in kernel panic.

    Fixes: 94cdb8bb993a ("cxgb4: Add support for dynamic allocation of resources for ULD")
    Signed-off-by: Shahjada Abul Husain <shahjada@chelsio.com>
    Signed-off-by: Vishal Kulkarni <vishal@chelsio.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d81be24c959b72328c342f187273ec3e4d261f75
Author: Alexandre Ghiti <alex@ghiti.fr>
Date:   Mon Sep 23 15:38:57 2019 -0700

    arm: use STACK_TOP when computing mmap base address

    [ Upstream commit 86e568e9c0525fc40e76d827212d5e9721cf7504 ]

    mmap base address must be computed wrt stack top address, using TASK_SIZE
    is wrong since STACK_TOP and TASK_SIZE are not equivalent.

    Link: http://lkml.kernel.org/r/20190730055113.23635-8-alex@ghiti.fr
    Signed-off-by: Alexandre Ghiti <alex@ghiti.fr>
    Acked-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Luis Chamberlain <mcgrof@kernel.org>
    Cc: Albert Ou <aou@eecs.berkeley.edu>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Christoph Hellwig <hch@infradead.org>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Palmer Dabbelt <palmer@sifive.com>
    Cc: Paul Burton <paul.burton@mips.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Will Deacon <will.deacon@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 53c12158540fea0b4833103a921346b2fed2708a
Author: Alexandre Ghiti <alex@ghiti.fr>
Date:   Mon Sep 23 15:38:54 2019 -0700
```

```
    arm: properly account for stack randomization and stack guard gap

    [ Upstream commit af0f4297286f13a75edf93677b1fb2fc16c412a7 ]

    This commit takes care of stack randomization and stack guard gap when
    computing mmap base address and checks if the task asked for
    randomization.  This fixes the problem uncovered and not fixed for arm
    here: https://lkml.kernel.org/r/20170622200033.25714-1-riel@redhat.com

    Link: http://lkml.kernel.org/r/20190730055113.23635-7-alex@ghiti.fr
    Signed-off-by: Alexandre Ghiti <alex@ghiti.fr>
    Acked-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Luis Chamberlain <mcgrof@kernel.org>
    Cc: Albert Ou <aou@eecs.berkeley.edu>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Christoph Hellwig <hch@infradead.org>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Palmer Dabbelt <palmer@sifive.com>
    Cc: Paul Burton <paul.burton@mips.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Will Deacon <will.deacon@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00645d66b56a1a946742f6d0e7c96be7db5fecef
Author: Alexandre Ghiti <alex@ghiti.fr>
Date:   Mon Sep 23 15:39:04 2019 -0700

    mips: properly account for stack randomization and stack guard gap

    [ Upstream commit b1f61b5bde3a1f50392c97b4c8513d1b8efb1cf2 ]

    This commit takes care of stack randomization and stack guard gap when
    computing mmap base address and checks if the task asked for
    randomization.  This fixes the problem uncovered and not fixed for arm
    here: https://lkml.kernel.org/r/20170622200033.25714-1-riel@redhat.com

    Link: http://lkml.kernel.org/r/20190730055113.23635-10-alex@ghiti.fr
    Signed-off-by: Alexandre Ghiti <alex@ghiti.fr>
    Acked-by: Kees Cook <keescook@chromium.org>
    Acked-by: Paul Burton <paul.burton@mips.com>
    Reviewed-by: Luis Chamberlain <mcgrof@kernel.org>
    Cc: Albert Ou <aou@eecs.berkeley.edu>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Christoph Hellwig <hch@infradead.org>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Palmer Dabbelt <palmer@sifive.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Will Deacon <will.deacon@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14dc3b6eee9732e76c7670c5fa5b73597b07b9a3
Author: Alexandre Ghiti <alex@ghiti.fr>
Date:   Mon Sep 23 15:38:43 2019 -0700
```

```
    arm64: consider stack randomization for mmap base only when necessary

    [ Upstream commit e8d54b62c55ab6201de6d195fc2c276294c1f6ae ]

    Do not offset mmap base address because of stack randomization if current
    task does not want randomization.  Note that x86 already implements this
    behaviour.

    Link: http://lkml.kernel.org/r/20190730055113.23635-4-alex@ghiti.fr
    Signed-off-by: Alexandre Ghiti <alex@ghiti.fr>
    Acked-by: Catalin Marinas <catalin.marinas@arm.com>
    Acked-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Luis Chamberlain <mcgrof@kernel.org>
    Cc: Albert Ou <aou@eecs.berkeley.edu>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Christoph Hellwig <hch@infradead.org>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Palmer Dabbelt <palmer@sifive.com>
    Cc: Paul Burton <paul.burton@mips.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Will Deacon <will.deacon@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b361c4dc338f03b1d810edcc6eec1edcb3f4579
Author: Nicolas Boichat <drinkcat@chromium.org>
Date:   Mon Sep 23 15:33:55 2019 -0700

    kmemleak: increase DEBUG_KMEMLEAK_EARLY_LOG_SIZE default to 16K

    [ Upstream commit b751c52bb587ae66f773b15204ef7a147467f4c7 ]

    The current default value (400) is too low on many systems (e.g.  some
    ARM64 platform takes up 1000+ entries).

    syzbot uses 16000 as default value, and has proved to be enough on beefy
    configurations, so let's pick that value.

    This consumes more RAM on boot (each entry is 160 bytes, so in total
    ~2.5MB of RAM), but the memory would later be freed (early_log is
    __initdata).

    Link: http://lkml.kernel.org/r/20190730154027.101525-1-drinkcat@chromium.org
    Signed-off-by: Nicolas Boichat <drinkcat@chromium.org>
    Suggested-by: Dmitry Vyukov <dvyukov@google.com>
    Acked-by: Catalin Marinas <catalin.marinas@arm.com>
    Acked-by: Dmitry Vyukov <dvyukov@google.com>
    Cc: Masahiro Yamada <yamada.masahiro@socionext.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Petr Mladek <pmladek@suse.com>
    Cc: Thomas Gleixner <tglx@linutronix.de>
    Cc: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Cc: Joe Lawrence <joe.lawrence@redhat.com>
    Cc: Uladzislau Rezki <urezki@gmail.com>
    Cc: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Cc: Stephen Rothwell <sfr@canb.auug.org.au>
    Cc: Andrey Ryabinin <aryabinin@virtuozzo.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad3cb1d921781475d93c19a1f7423f4e9afd6737
```

```
Author: Changwei Ge <gechangwei@live.cn>
Date:   Mon Sep 23 15:33:37 2019 -0700

    ocfs2: wait for recovering done after direct unlock request

    [ Upstream commit 0a3775e4f883912944481cf2ef36eb6383a9cc74 ]

    There is a scenario causing ocfs2 umount hang when multiple hosts are
    rebooting at the same time.

    NODE1                           NODE2               NODE3
    send unlock requset to NODE2
                                    dies
                                                        become recovery master
                                                        recover NODE2
    find NODE2 dead
    mark resource RECOVERING
    directly remove lock from grant list
    calculate usage but RECOVERING marked
    **miss the window of purging
    clear RECOVERING

    To reproduce this issue, crash a host and then umount ocfs2
    from another node.

    To solve this, just let unlock progress wait for recovery done.

    Link: http://lkml.kernel.org/r/1550124866-20367-1-git-send-email-gechangwei@live.cn
    Signed-off-by: Changwei Ge <gechangwei@live.cn>
    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>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ce7a207061dbde29a33093705ea0fe21a4308b9
Author: Greg Thelen <gthelen@google.com>
Date:   Mon Sep 23 15:33:05 2019 -0700

    kbuild: clean compressed initramfs image

    [ Upstream commit 6279eb3dd7946c69346a3b98473ed13d3a44adb5 ]

    Since 9e3596b0c653 ("kbuild: initramfs cleanup, set target from Kconfig")
    "make clean" leaves behind compressed initramfs images.  Example:

      $ make defconfig
      $ sed -i 's|CONFIG_INITRAMFS_SOURCE=""|CONFIG_INITRAMFS_SOURCE="/tmp/ir.cpio"|'
.config
      $ make olddefconfig
      $ make -s
      $ make -s clean
      $ git clean -ndxf | grep initramfs
      Would remove usr/initramfs_data.cpio.gz

    clean rules do not have CONFIG_* context so they do not know which
    compression format was used.  Thus they don't know which files to delete.

    Tell clean to delete all possible compression formats.

    Once patched usr/initramfs_data.cpio.gz and friends are deleted by
    "make clean".
```

```
    Link: http://lkml.kernel.org/r/20190722063251.55541-1-gthelen@google.com
    Fixes: 9e3596b0c653 ("kbuild: initramfs cleanup, set target from Kconfig")
    Signed-off-by: Greg Thelen <gthelen@google.com>
    Cc: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac246294131ba6389db7b492fb4c6b0696964c4c
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Sep 23 15:32:59 2019 -0700

    mm: add dummy can_do_mlock() helper

    [ Upstream commit 710ec38b0f633ab3e2581f07a73442d809e28ab0 ]

    On kernels without CONFIG_MMU, we get a link error for the siw driver:

    drivers/infiniband/sw/siw/siw_mem.o: In function `siw_umem_get':
    siw_mem.c:(.text+0x4c8): undefined reference to `can_do_mlock'

    This is probably not the only driver that needs the function and could
    otherwise build correctly without CONFIG_MMU, so add a dummy variant that
    always returns false.

    Link: http://lkml.kernel.org/r/20190909204201.931830-1-arnd@arndb.de
    Fixes: 2251334dcac9 ("rdma/siw: application buffer management")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Suggested-by: Jason Gunthorpe <jgg@mellanox.com>
    Acked-by: Michal Hocko <mhocko@suse.com>
    Cc: Bernard Metzler <bmt@zurich.ibm.com>
    Cc: "Matthew Wilcox (Oracle)" <willy@infradead.org>
    Cc: "Kirill A. Shutemov" <kirill.shutemov@linux.intel.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e9b0526f2635f4f563d5f0e2278569d4426d3fd
Author: Yunfeng Ye <yeyunfeng@huawei.com>
Date:   Sun Sep 15 17:26:56 2019 +0800

    crypto: hisilicon - Fix double free in sec_free_hw_sgl()

    [ Upstream commit 24fbf7bad888767bed952f540ac963bc57e47e15 ]

    There are two problems in sec_free_hw_sgl():

    First, when sgl_current->next is valid, @hw_sgl will be freed in the
    first loop, but it free again after the loop.

    Second, sgl_current and sgl_current->next_sgl is not match when
    dma_pool_free() is invoked, the third parameter should be the dma
    address of sgl_current, but sgl_current->next_sgl is the dma address
    of next chain, so use sgl_current->next_sgl is wrong.

    Fix this by deleting the last dma_pool_free() in sec_free_hw_sgl(),
    modifying the condition for while loop, and matching the address for
    dma_pool_free().

    Fixes: 915e4e8413da ("crypto: hisilicon - SEC security accelerator driver")
    Signed-off-by: Yunfeng Ye <yeyunfeng@huawei.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 547fcb77abab04df66dd13892e817e593db5a548
Author: Youquan Song <youquan.song@intel.com>
Date:   Sat Sep 14 12:45:43 2019 -0700

    tools/power/x86/intel-speed-select: Fix high priority core mask over count

    [ Upstream commit 44460efe44e05eae2f21e57d06d542bbbb792e65 ]

    If the CPU package has the less logical CPU than topo_max_cpus, but un-present
    CPU's punit_cpu_core will be initiated to 0 and they will be count to core 0

    Like below, there are only 10 high priority cores (20 logical CPUs) in the CPU
    package, but it count to 27 logic CPUs.

    ./intel-speed-select base-freq info -l 0 | grep mask
            high-priority-cpu-mask:7f000179,f000179f

    With the fix patch:
    ./intel-speed-select base-freq info -l 0
            high-priority-cpu-mask:00000179,f000179f

    Signed-off-by: Youquan Song <youquan.song@intel.com>
    Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6270457c4aa2d25840d2ef18f2f2386bd24dd3da
Author: David Howells <dhowells@redhat.com>
Date:   Thu Mar 21 10:08:08 2019 +0000

    hypfs: Fix error number left in struct pointer member

    [ Upstream commit b54c64f7adeb241423cd46598f458b5486b0375e ]

    In hypfs_fill_super(), if hypfs_create_update_file() fails,
    sbi->update_file is left holding an error number.  This is passed to
    hypfs_kill_super() which doesn't check for this.

    Fix this by not setting sbi->update_value until after we've checked for
    error.

    Fixes: 24bbb1faf3f0 ("[PATCH] s390_hypfs filesystem")
    Signed-off-by: David Howells <dhowells@redhat.com>
    cc: Martin Schwidefsky <schwidefsky@de.ibm.com>
    cc: Heiko Carstens <heiko.carstens@de.ibm.com>
    cc: linux-s390@vger.kernel.org
    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c8fc5b0209c157a4ab0ba41edd35e0d2335db7ba
Author: Jens Axboe <axboe@kernel.dk>
Date:   Sun Sep 22 10:01:05 2019 -0600

    pktcdvd: remove warning on attempting to register non-passthrough dev

    [ Upstream commit eb09b3cc464d2c3bbde9a6648603c8d599ea8582 ]

    Anatoly reports that he gets the below warning when booting -git on
    a sparc64 box on debian unstable:


    ...
    [   13.352975] aes_sparc64: Using sparc64 aes opcodes optimized AES
    implementation
    [   13.428002] ------------[ cut here ]------------
    [   13.428081] WARNING: CPU: 21 PID: 586 at
```

```
    drivers/block/pktcdvd.c:2597 pkt_setup_dev+0x2e4/0x5a0 [pktcdvd]
    [   13.428147] Attempt to register a non-SCSI queue
    [   13.428184] Modules linked in: pktcdvd libdes cdrom aes_sparc64
    n2_rng md5_sparc64 sha512_sparc64 rng_core sha256_sparc64 flash
    sha1_sparc64 ip_tables x_tables ipv6 crc_ccitt nf_defrag_ipv6 autofs4
    ext4 crc16 mbcache jbd2 raid10 raid456 async_raid6_recov async_memcpy
    async_pq async_xor xor async_tx raid6_pq raid1 raid0 multipath linear
    md_mod crc32c_sparc64
    [   13.428452] CPU: 21 PID: 586 Comm: pktsetup Not tainted
    5.3.0-10169-g574cc4539762 #1234
    [   13.428507] Call Trace:
    [   13.428542]  [0000000004635c0] __warn+0xc0/0x100
    [   13.428582]  [0000000000463634] warn_slowpath_fmt+0x34/0x60
    [   13.428626]  [000000001045b244] pkt_setup_dev+0x2e4/0x5a0 [pktcdvd]
    [   13.428674]  [000000001045ccf4] pkt_ctl_ioctl+0x94/0x220 [pktcdvd]
    [   13.428724]  [00000000006b95c8] do_vfs_ioctl+0x628/0x6e0
    [   13.428764]  [00000000006b96c8] ksys_ioctl+0x48/0x80
    [   13.428803]  [00000000006b9714] sys_ioctl+0x14/0x40
    [   13.428847]  [0000000000406294] linux_sparc_syscall+0x34/0x44
    [   13.428890] irq event stamp: 4181
    [   13.428924] hardirqs last  enabled at (4189): [<00000000004e0a74>]
    console_unlock+0x634/0x6c0
    [   13.428984] hardirqs last disabled at (4196): [<00000000004e0540>]
    console_unlock+0x100/0x6c0
    [   13.429048] softirqs last  enabled at (3978): [<0000000000b2e2d8>]
    __do_softirq+0x498/0x520
    [   13.429110] softirqs last disabled at (3967): [<000000000042cfb4>]
    do_softirq_own_stack+0x34/0x60
    [   13.429172] ---[ end trace 2220ca468f32967d ]---
    [   13.430018] pktcdvd: setup of pktcdvd device failed
    [   13.455589] des_sparc64: Using sparc64 des opcodes optimized DES
    implementation
    [   13.515334] camellia_sparc64: Using sparc64 camellia opcodes
    optimized CAMELLIA implementation
    [   13.522856] pktcdvd: setup of pktcdvd device failed
    [   13.529327] pktcdvd: setup of pktcdvd device failed
    [   13.532932] pktcdvd: setup of pktcdvd device failed
    [   13.536165] pktcdvd: setup of pktcdvd device failed
    [   13.539372] pktcdvd: setup of pktcdvd device failed
    [   13.542834] pktcdvd: setup of pktcdvd device failed
    [   13.546536] pktcdvd: setup of pktcdvd device failed
    [   15.431071] XFS (dm-0): Mounting V5 Filesystem
    ...

    Apparently debian auto-attaches any cdrom like device to pktcdvd, which
    can lead to the above warning. There's really no reason to warn for this
    situation, kill it.

    Reported-by: Anatoly Pugachev <matorola@gmail.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5044962d0976b527b817e58ebd0ac5bed0d0b68e
Author: OGAWA Hirofumi <hirofumi@mail.parknet.co.jp>
Date:   Mon Sep 23 15:32:53 2019 -0700

    fat: work around race with userspace's read via blockdev while mounting

    [ Upstream commit 07bfa4415ab607e459b69bd86aa7e7602ce10b4f ]

    If userspace reads the buffer via blockdev while mounting,
    sb_getblk()+modify can race with buffer read via blockdev.

    For example,
```

```
                FS                                  userspace
        bh = sb_getblk()
        modify bh->b_data
                                      read
                                        ll_rw_block(bh)
                                          fill bh->b_data by on-disk data
                                          /* lost modified data by FS */
                                          set_buffer_uptodate(bh)
        set_buffer_uptodate(bh)
```

    Userspace should not use the blockdev while mounting though, the udev
    seems to be already doing this. Although I think the udev should try to
    avoid this, workaround the race by small overhead.

    Link: http://lkml.kernel.org/r/87pnk7l3sw.fsf_-_@mail.parknet.co.jp
    Signed-off-by: OGAWA Hirofumi <hirofumi@mail.parknet.co.jp>
    Reported-by: Jan Stancek <jstancek@redhat.com>
    Tested-by: Jan Stancek <jstancek@redhat.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8228d4ed13d145cbf02bf262910aecfcd9a25b8
Author: Paolo Valente <paolo.valente@linaro.org>
Date: Thu Aug 22 17:20:37 2019 +0200

    block, bfq: push up injection only after setting service time

    [Upstream commit 58494c980f40274c465ebfdece02d401def088bf]

    If equal to 0, the injection limit for a bfq_queue is pushed to 1
    after a first sample of the total service time of the I/O requests of
    the queue is computed (to allow injection to start). Yet, because of a
    mistake in the branch that performs this action, the push may happen
    also in some other case. This commit fixes this issue.

    Tested-by: Oleksandr Natalenko <oleksandr@natalenko.name>
    Signed-off-by: Paolo Valente <paolo.valente@linaro.org>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de8d059ffbd669ca00311548b8851341eb89ee5f
Author: Jon Hunter <jonathanh@nvidia.com>
Date: Tue Sep 10 10:29:17 2019 +0100

    i2c: tegra: Move suspend handling to NOIRQ phase

    [Upstream commit 8ebf15e9c869e764b3aab4928938ee68c0e2bd6d]

    Commit acc8abcb2a9c ("i2c: tegra: Add suspend-resume support") added
    suspend support for the Tegra I2C driver and following this change on
    Tegra30 the following WARNING is seen on entering suspend ...

     WARNING: CPU: 2 PID: 689 at /dvs/git/dirty/git-master_l4t-
upstream/kernel/drivers/i2c/i2c-core.h:54 __i2c_transfer+0x35c/0x70c
     i2c i2c-4: Transfer while suspended
     Modules linked in: brcmfmac brcmutil
     CPU: 2 PID: 689 Comm: rtcwake Not tainted 5.3.0-rc7-g089cf7f6ecb2 #1
     Hardware name: NVIDIA Tegra SoC (Flattened Device Tree)
     [<c0112264>] (unwind_backtrace) from [<c010ca94>] (show_stack+0x10/0x14)
     [<c010ca94>] (show_stack) from [<c0a77024>] (dump_stack+0xb4/0xc8)
     [<c0a77024>] (dump_stack) from [<c0124198>] (__warn+0xe0/0xf8)
     [<c0124198>] (__warn) from [<c01241f8>] (warn_slowpath_fmt+0x48/0x6c)
     [<c01241f8>] (warn_slowpath_fmt) from [<c06f6c40>] (__i2c_transfer+0x35c/0x70c)
     [<c06f6c40>] (__i2c_transfer) from [<c06f7048>] (i2c_transfer+0x58/0xf4)

```
      [<c06f7048>] (i2c_transfer) from [<c06f7130>] (i2c_transfer_buffer_flags+0x4c/0x70)
      [<c06f7130>] (i2c_transfer_buffer_flags) from [<c05bee78>] (regmap_i2c_write+0x14/0x30)
      [<c05bee78>] (regmap_i2c_write) from [<c05b9cac>] (_regmap_raw_write_impl+0x35c/0x868)
      [<c05b9cac>] (_regmap_raw_write_impl) from [<c05b984c>] (_regmap_update_bits+0xe4/0xec)
      [<c05b984c>] (_regmap_update_bits) from [<c05bad04>]
  (regmap_update_bits_base+0x50/0x74)
      [<c05bad04>] (regmap_update_bits_base) from [<c04d453c>]
  (regulator_disable_regmap+0x44/0x54)
      [<c04d453c>] (regulator_disable_regmap) from [<c04cf9d4>]
  (_regulator_do_disable+0xf8/0x268)
      [<c04cf9d4>] (_regulator_do_disable) from [<c04d1694>] (_regulator_disable+0xf4/0x19c)
      [<c04d1694>] (_regulator_disable) from [<c04d1770>] (regulator_disable+0x34/0x64)
      [<c04d1770>] (regulator_disable) from [<c04d2310>] (regulator_bulk_disable+0x28/0xb4)
      [<c04d2310>] (regulator_bulk_disable) from [<c0495cd4>]
  (tegra_pcie_power_off+0x64/0xa8)
      [<c0495cd4>] (tegra_pcie_power_off) from [<c0495f74>]
  (tegra_pcie_pm_suspend+0x25c/0x3f4)
      [<c0495f74>] (tegra_pcie_pm_suspend) from [<c05af48c>] (dpm_run_callback+0x38/0x1d4)
      [<c05af48c>] (dpm_run_callback) from [<c05afe30>] (__device_suspend_noirq+0xc0/0x2b8)
      [<c05afe30>] (__device_suspend_noirq) from [<c05b1c24>]
  (dpm_noirq_suspend_devices+0x100/0x37c)
      [<c05b1c24>] (dpm_noirq_suspend_devices) from [<c05b1ebc>]
  (dpm_suspend_noirq+0x1c/0x48)
      [<c05b1ebc>] (dpm_suspend_noirq) from [<c017d2c0>]
  (suspend_devices_and_enter+0x1d0/0xa00)
      [<c017d2c0>] (suspend_devices_and_enter) from [<c017dd10>] (pm_suspend+0x220/0x74c)
      [<c017dd10>] (pm_suspend) from [<c017c2c8>] (state_store+0x6c/0xc8)
      [<c017c2c8>] (state_store) from [<c02ef398>] (kernfs_fop_write+0xe8/0x1c4)
      [<c02ef398>] (kernfs_fop_write) from [<c0271e38>] (__vfs_write+0x2c/0x1c4)
      [<c0271e38>] (__vfs_write) from [<c02748dc>] (vfs_write+0xa4/0x184)
      [<c02748dc>] (vfs_write) from [<c0274b7c>] (ksys_write+0x9c/0xdc)
      [<c0274b7c>] (ksys_write) from [<c0101000>] (ret_fast_syscall+0x0/0x54)
     Exception stack(0xe9f21fa8 to 0xe9f21ff0)
      1fa0:                   0000006c 004b2438 00000004 004b2438 00000004 00000000
      1fc0: 0000006c 004b2438 004b1228 00000004 00000004 00000004 0049e78c 004b1228
      1fe0: 00000004 be9809b8 b6f0bc0b b6e96206

    The problem is that the Tegra PCIe driver indirectly uses I2C for
    controlling some regulators and the I2C driver is now being suspended
    before the PCIe driver causing the PCIe suspend to fail. The Tegra PCIe
    driver is suspended during the NOIRQ phase and this cannot be changed
    due to other dependencies. Therefore, we also need to move the suspend
    handling for the Tegra I2C driver to the NOIRQ phase as well.

    In order to move the I2C suspend handling to the NOIRQ phase we also
    need to avoid calling pm_runtime_get/put() because per commit
    1e2ef05bb8cf ("PM: Limit race conditions between runtime PM and system
    sleep (v2)") these cannot be called early in resume. The function
    tegra_i2c_init(), called during resume, calls pm_runtime_get/put() and
    so move these calls outside of tegra_i2c_init(), so this function can
    be used during the NOIRQ resume phase.

    Fixes: acc8abcb2a9c ("i2c: tegra: Add suspend-resume support")

    Signed-off-by: Jon Hunter <jonathanh@nvidia.com>
    Acked-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Wolfram Sang <wsa@the-dreams.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f9affab8498d0f24a9f767711b2f22dc4bcff781
Author: Mike Rapoport <mike.rapoport@gmail.com>
Date:   Fri Aug 30 14:27:56 2019 +0100

    ARM: 8903/1: ensure that usable memory in bank 0 starts from a PMD-aligned address
```

```
    [ Upstream commit 00d2ec1e6bd82c0538e6dd3e4a4040de93ba4fef ]

    The calculation of memblock_limit in adjust_lowmem_bounds() assumes that
    bank 0 starts from a PMD-aligned address. However, the beginning of the
    first bank may be NOMAP memory and the start of usable memory
    will be not aligned to PMD boundary. In such case the memblock_limit will
    be set to the end of the NOMAP region, which will prevent any memblock
    allocations.

    Mark the region between the end of the NOMAP area and the next PMD-aligned
    address as NOMAP as well, so that the usable memory will start at
    PMD-aligned address.

    Signed-off-by: Mike Rapoport <rppt@linux.ibm.com>
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4811149cbcc1ea381ee6487aebf552c227ce0d3d
Author: Nathan Chancellor <natechancellor@gmail.com>
Date:   Wed Sep 4 01:13:15 2019 +0100

    ARM: 8905/1: Emit __gnu_mcount_nc when using Clang 10.0.0 or newer

    [ Upstream commit b0fe66cf095016e0b238374c10ae366e1f087d11 ]

    Currently, multi_v7_defconfig + CONFIG_FUNCTION_TRACER fails to build
    with clang:

    arm-linux-gnueabi-ld: kernel/softirq.o: in function `_local_bh_enable':
    softirq.c:(.text+0x504): undefined reference to `mcount'
    arm-linux-gnueabi-ld: kernel/softirq.o: in function `__local_bh_enable_ip':
    softirq.c:(.text+0x58c): undefined reference to `mcount'
    arm-linux-gnueabi-ld: kernel/softirq.o: in function `do_softirq':
    softirq.c:(.text+0x6c8): undefined reference to `mcount'
    arm-linux-gnueabi-ld: kernel/softirq.o: in function `irq_enter':
    softirq.c:(.text+0x75c): undefined reference to `mcount'
    arm-linux-gnueabi-ld: kernel/softirq.o: in function `irq_exit':
    softirq.c:(.text+0x840): undefined reference to `mcount'
    arm-linux-gnueabi-ld: kernel/softirq.o:softirq.c:(.text+0xa50): more undefined
references to `mcount' follow

    clang can emit a working mcount symbol, __gnu_mcount_nc, when
    '-meabi gnu' is passed to it. Until r369147 in LLVM, this was
    broken and caused the kernel not to boot with '-pg' because the
    calling convention was not correct. Always build with '-meabi gnu'
    when using clang but ensure that '-pg' (which is added with
    CONFIG_FUNCTION_TRACER and its prereq CONFIG_HAVE_FUNCTION_TRACER)
    cannot be added with it unless this is fixed (which means using
    clang 10.0.0 and newer).

    Link: https://github.com/ClangBuiltLinux/linux/issues/35
    Link: https://bugs.llvm.org/show_bug.cgi?id=33845
    Link: https://github.com/llvm/llvm-
project/commit/16fa8b09702378bacfa3d07081afe6b353b99e60

    Reviewed-by: Matthias Kaehlcke <mka@chromium.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Reviewed-by: Stefan Agner <stefan@agner.ch>
    Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e14f6865487eedcd2ac68034e17e9e6b17f8f18c
Author: Krzysztof Wilczynski <kw@linux.com>
Date:   Mon Aug 26 17:14:36 2019 +0200
```

```
    PCI: Use static const struct, not const static struct

    [ Upstream commit 8050f3f6645ae0f7e4c1304593f6f7eb2ee7d85c ]

    Move the static keyword to the front of declarations of pci_regs_behavior[]
    and pcie_cap_regs_behavior[], which resolves compiler warnings when
    building with "W=1":

      drivers/pci/pci-bridge-emul.c:41:1: warning: ‘static’ is not at beginning of
      declaration [-Wold-style-declaration]
       const static struct pci_bridge_reg_behavior pci_regs_behavior[] = {
       ^
      drivers/pci/pci-bridge-emul.c:176:1: warning: ‘static’ is not at beginning of
      declaration [-Wold-style-declaration]
       const static struct pci_bridge_reg_behavior pcie_cap_regs_behavior[] = {
       ^

    Link: https://lore.kernel.org/r/20190826151436.4672-1-kw@linux.com
    Link: https://lore.kernel.org/r/20190828131733.5817-1-kw@linux.com
    Signed-off-by: Krzysztof Wilczynski <kw@linux.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Thomas Petazzoni <thomas.petazzoni@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 397a607551d485877f2218fcd8e077d0a6c2e1eb
Author: Jia-Ju Bai <baijiaju1990@gmail.com>
Date:   Tue Jul 23 18:00:15 2019 +0800

    security: smack: Fix possible null-pointer dereferences in smack_socket_sock_rcv_skb()

    [ Upstream commit 3f4287e7d98a2954f20bf96c567fdffcd2b63eb9 ]

    In smack_socket_sock_rcv_skb(), there is an if statement
    on line 3920 to check whether skb is NULL:
        if (skb && skb->secmark != 0)

    This check indicates skb can be NULL in some cases.

    But on lines 3931 and 3932, skb is used:
        ad.a.u.net->netif = skb->skb_iif;
        ipv6_skb_to_auditdata(skb, &ad.a, NULL);

    Thus, possible null-pointer dereferences may occur when skb is NULL.

    To fix these possible bugs, an if statement is added to check skb.

    These bugs are found by a static analysis tool STCheck written by us.

    Signed-off-by: Jia-Ju Bai <baijiaju1990@gmail.com>
    Signed-off-by: Casey Schaufler <casey@schaufler-ca.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a4e61e26aa9dd9e4eba43c21c0b862d3ec786b65
Author: Thierry Reding <treding@nvidia.com>
Date:   Thu Aug 29 12:53:15 2019 +0200

    PCI: exynos: Propagate errors for optional PHYs

    [ Upstream commit ddd6960087d4b45759434146d681a94bbb1c54ad ]

    devm_of_phy_get() can fail for a number of reasons besides probe
    deferral. It can for example return -ENOMEM if it runs out of memory as
    it tries to allocate devres structures. Propagating only -EPROBE_DEFER
    is problematic because it results in these legitimately fatal errors
```

```
    being treated as "PHY not specified in DT".

    What we really want is to ignore the optional PHYs only if they have not
    been specified in DT. devm_of_phy_get() returns -ENODEV in this case, so
    that's the special case that we need to handle. So we propagate all
    errors, except -ENODEV, so that real failures will still cause the
    driver to fail probe.

    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Andrew Murray <andrew.murray@arm.com>
    Cc: Jingoo Han <jingoohan1@gmail.com>
    Cc: Kukjin Kim <kgene@kernel.org>
    Cc: Krzysztof Kozlowski <krzk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4b8763f7deaaf800002087de3c5c66d627df59ec
Author: Thierry Reding <treding@nvidia.com>
Date:   Thu Aug 29 12:53:16 2019 +0200

    PCI: imx6: Propagate errors for optional regulators

    [ Upstream commit 2170a09fb4b0f66e06e5bcdcbc98c9ccbf353650 ]

    regulator_get_optional() can fail for a number of reasons besides probe
    deferral. It can for example return -ENOMEM if it runs out of memory as
    it tries to allocate data structures. Propagating only -EPROBE_DEFER is
    problematic because it results in these legitimately fatal errors being
    treated as "regulator not specified in DT".

    What we really want is to ignore the optional regulators only if they
    have not been specified in DT. regulator_get_optional() returns -ENODEV
    in this case, so that's the special case that we need to handle. So we
    propagate all errors, except -ENODEV, so that real failures will still
    cause the driver to fail probe.

    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Andrew Murray <andrew.murray@arm.com>
    Cc: Richard Zhu <hongxing.zhu@nxp.com>
    Cc: Lucas Stach <l.stach@pengutronix.de>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Cc: Sascha Hauer <s.hauer@pengutronix.de>
    Cc: Fabio Estevam <festevam@gmail.com>
    Cc: kernel@pengutronix.de
    Cc: linux-imx@nxp.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2caacb92a87fcf97abb9f9743157b8dcf09c84d3
Author: Thierry Reding <treding@nvidia.com>
Date:   Thu Aug 29 12:53:18 2019 +0200

    PCI: histb: Propagate errors for optional regulators

    [ Upstream commit 8f9e1641ba445437095411d9fda2324121110d5d ]

    regulator_get_optional() can fail for a number of reasons besides probe
    deferral. It can for example return -ENOMEM if it runs out of memory as
    it tries to allocate data structures. Propagating only -EPROBE_DEFER is
    problematic because it results in these legitimately fatal errors being
    treated as "regulator not specified in DT".

    What we really want is to ignore the optional regulators only if they
    have not been specified in DT. regulator_get_optional() returns -ENODEV
    in this case, so that's the special case that we need to handle. So we
```

```
    propagate all errors, except -ENODEV, so that real failures will still
    cause the driver to fail probe.

    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Andrew Murray <andrew.murray@arm.com>
    Cc: Shawn Guo <shawn.guo@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2a0976536241a1e7e259eda1850bd049214e1ee7
Author: Thierry Reding <treding@nvidia.com>
Date:   Thu Aug 29 12:53:14 2019 +0200

    PCI: rockchip: Propagate errors for optional regulators

    [ Upstream commit 0e3ff0ac5f71bdb6be2a698de0ed0c7e6e738269 ]

    regulator_get_optional() can fail for a number of reasons besides probe
    deferral. It can for example return -ENOMEM if it runs out of memory as
    it tries to allocate data structures. Propagating only -EPROBE_DEFER is
    problematic because it results in these legitimately fatal errors being
    treated as "regulator not specified in DT".

    What we really want is to ignore the optional regulators only if they
    have not been specified in DT. regulator_get_optional() returns -ENODEV
    in this case, so that's the special case that we need to handle. So we
    propagate all errors, except -ENODEV, so that real failures will still
    cause the driver to fail probe.

    Tested-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Andrew Murray <andrew.murray@arm.com>
    Reviewed-by: Heiko Stuebner <heiko@sntech.de>
    Acked-by: Shawn Lin <shawn.lin@rock-chips.com>
    Cc: Shawn Lin <shawn.lin@rock-chips.com>
    Cc: Heiko Stuebner <heiko@sntech.de>
    Cc: linux-rockchip@lists.infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89f817442e65b7eb9aae80c8b5fec2dafc49ae25
Author: Joao Moreno <mail@joaomoreno.com>
Date:   Tue Sep 3 16:46:32 2019 +0200

    HID: apple: Fix stuck function keys when using FN

    [ Upstream commit aec256d0ecd561036f188dbc8fa7924c47a9edfd ]

    This fixes an issue in which key down events for function keys would be
    repeatedly emitted even after the user has raised the physical key. For
    example, the driver fails to emit the F5 key up event when going through
    the following steps:
    - fnmode=1: hold FN, hold F5, release FN, release F5
    - fnmode=2: hold F5, hold FN, release F5, release FN

    The repeated F5 key down events can be easily verified using xev.

    Signed-off-by: Joao Moreno <mail@joaomoreno.com>
    Co-developed-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e1c023a877143dea96fef5f0b973c69326609b4
Author: Krzysztof Wilczynski <kw@linux.com>
Date:   Mon Aug 26 00:46:16 2019 +0200
```

```
    PCI: Add pci_info_ratelimited() to ratelimit PCI separately

    [ Upstream commit 7f1c62c443a453deb6eb3515e3c05650ffe0dcf0 ]

    Do not use printk_ratelimit() in drivers/pci/pci.c as it shares the rate
    limiting state with all other callers to the printk_ratelimit().

    Add pci_info_ratelimited() (similar to pci_notice_ratelimited() added in
    the commit a88a7b3eb076 ("vfio: Use dev_printk() when possible")) and use
    it instead of printk_ratelimit() + pci_info().

    Link: https://lore.kernel.org/r/20190825224616.8021-1-kw@linux.com
    Signed-off-by: Krzysztof Wilczynski <kw@linux.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2af7fbcc79816cc1013d40782654f3cde4ac8c4f
Author: Stephen Smalley <sds@tycho.nsa.gov>
Date:   Wed Sep 4 10:32:48 2019 -0400

    selinux: fix residual uses of current_security() for the SELinux blob

    [ Upstream commit 169ce0c081cd85f78388bb6c1638c1ad7b81bde7 ]

    We need to use selinux_cred() to fetch the SELinux cred blob instead
    of directly using current->security or current_security().  There
    were a couple of lingering uses of current_security() in the SELinux code
    that were apparently missed during the earlier conversions. IIUC, this
    would only manifest as a bug if multiple security modules including
    SELinux are enabled and SELinux is not first in the lsm order. After
    this change, there appear to be no other users of current_security()
    in-tree; perhaps we should remove it altogether.

    Fixes: bbd3662a8348 ("Infrastructure management of the cred security blob")
    Signed-off-by: Stephen Smalley <sds@tycho.nsa.gov>
    Acked-by: Casey Schaufler <casey@schaufler-ca.com>
    Reviewed-by: James Morris <jamorris@linux.microsoft.com>
    Signed-off-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d076d0facf7bc24469501fc9d29e38338df6ad70
Author: Romain Izard <romain.izard.pro@gmail.com>
Date:   Fri Aug 30 15:15:56 2019 +0200

    power: supply: register HWMON devices with valid names

    [ Upstream commit f1b937cc86bedf94dbc3478c2c0dc79471081fff ]

    With the introduction of the HWMON compatibility layer to the power
    supply framework in Linux 5.3, all power supply devices' names can be
    used directly to create HWMON devices with the same names.

    But HWMON has rules on allowable names that are different from those
    used in the power supply framework. The dash character is forbidden, as
    it is used by the libsensors library in userspace as a separator,
    whereas this character is used in the device names in more than half of
    the existing power supply drivers. This last case is consistent with the
    typical naming usage with MFD and Device Tree.

    This leads to warnings in the kernel log, with the format:

    power_supply gpio-charger: hwmon: \
            'gpio-charger' is not a valid name attribute, please fix
```

```
    Add a protection to power_supply_add_hwmon_sysfs() that replaces any
    dash in the device name with an underscore when registering with the
    HWMON framework. Other forbidden characters (star, slash, space, tab,
    newline) are not replaced, as they are not in common use.

    Fixes: e67d4dfc9ff1 ("power: supply: Add HWMON compatibility layer")
    Signed-off-by: Romain Izard <romain.izard.pro@gmail.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7974bdbb45fac11d98cc52e58e500d0f6126de47
Author: Biwen Li <biwen.li@nxp.com>
Date:   Thu Aug 29 10:14:18 2019 +0800

    rtc: pcf85363/pcf85263: fix regmap error in set_time

    [ Upstream commit 7ef66122bdb3b839e9f51b76d7e600b6e21ef648 ]

    Issue:
        - # hwclock -w
          hwclock: RTC_SET_TIME: Invalid argument

    Why:
        - Relative commit: 8b9f9d4dc511 ("regmap: verify if register is
          writeable before writing operations"), this patch
          will always check for unwritable registers, it will compare reg
          with max_register in regmap_writeable.

        - The pcf85363/pcf85263 has the capability of address wrapping
          which means if you access an address outside the allowed range
          (0x00-0x2f) hardware actually wraps the access to a lower address.
          The rtc-pcf85363 driver will use this feature to configure the time
          and execute 2 actions in the same i2c write operation (stopping the
          clock and configure the time). However the driver has also
          configured the `regmap maxregister` protection mechanism that will
          block accessing addresses outside valid range (0x00-0x2f).

    How:
        - Split of writing regs to two parts, first part writes control
          registers about stop_enable and resets, second part writes
          RTC time and date registers.

    Signed-off-by: Biwen Li <biwen.li@nxp.com>
    Link: https://lore.kernel.org/r/20190829021418.4607-1-biwen.li@nxp.com
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4bc8b4d7acde9489eef72cbefefe627f454c79bd
Author: Anson Huang <Anson.Huang@nxp.com>
Date:   Tue Jul 16 15:18:58 2019 +0800

    rtc: snvs: fix possible race condition

    [ Upstream commit 6fd4fe9b496d9ba3382992ff4fde3871d1b6f63d ]

    The RTC IRQ is requested before the struct rtc_device is allocated,
    this may lead to a NULL pointer dereference in IRQ handler.

    To fix this issue, allocating the rtc_device struct before requesting
    the RTC IRQ using devm_rtc_allocate_device, and use rtc_register_device
    to register the RTC device.

    Signed-off-by: Anson Huang <Anson.Huang@nxp.com>
    Reviewed-by: Dong Aisheng <aisheng.dong@nxp.com>
```

```
    Link: https://lore.kernel.org/r/20190716071858.36750-1-Anson.Huang@nxp.com
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68fec9c993dc1eb86027388595417c3e7e293f03
Author: Nick Desaulniers <ndesaulniers@google.com>
Date:   Mon Jul 8 20:38:15 2019 +0100

    ARM: 8875/1: Kconfig: default to AEABI w/ Clang

    [ Upstream commit a05b9608456e0d4464c6f7ca8572324ace57a3f4 ]

    Clang produces references to __aeabi_uidivmod and __aeabi_idivmod for
    arm-linux-gnueabi and arm-linux-gnueabihf targets incorrectly when AEABI
    is not selected (such as when OABI_COMPAT is selected).

    While this means that OABI userspaces wont be able to upgraded to
    kernels built with Clang, it means that boards that don't enable AEABI
    like s3c2410_defconfig will stop failing to link in KernelCI when built
    with Clang.

    Link: https://github.com/ClangBuiltLinux/linux/issues/482
    Link: https://groups.google.com/forum/#!msg/clang-built-linux/yydsAAux5hk/GxjqJSW-AQAJ

    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Nick Desaulniers <ndesaulniers@google.com>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15c679adcfad9daaba4c71e7f76ef6e751e1bdbf
Author: Hou Zhiqiang <Zhiqiang.Hou@nxp.com>
Date:   Sat Jul 13 22:11:29 2019 +0800

    PCI: mobiveil: Fix the CPU base address setup in inbound window

    [ Upstream commit df901c85cc28b538c62f6bc20b16a8bd05fcb756 ]

    Current code erroneously sets-up the CPU base address through the
    parameter 'pci_addr', which is passed to initialize the CPU (AXI) base
    address of the inbound window where the controller maps the PCI address
    space into CPU physical address space; furthermore, it also truncates it
    by programming only the lower 32-bit value into the inbound CPU address
    register.

    Fix both issues by introducing a new parameter 'u64 cpu_addr' to
    initialize both lower 32-bit and upper 32-bit of the CPU physical
    base address mapping PCI inbound transactions into CPU (AXI) ones.

    Fixes: 9af6bcb11e12 ("PCI: mobiveil: Add Mobiveil PCIe Host Bridge IP driver")
    Signed-off-by: Hou Zhiqiang <Zhiqiang.Hou@nxp.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Minghuan Lian <Minghuan.Lian@nxp.com>
    Reviewed-by: Subrahmanya Lingappa <l.subrahmanya@mobiveil.co.in>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d101356c6ecc09a996bbaa344ad64e7cc54b4ee1
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Aug 5 19:55:07 2019 -0500

    soundwire: intel: fix channel number reported by hardware

    [ Upstream commit 18046335643de6d21327f5ae034c8fb8463f6715 ]
```

```
    On all released Intel controllers (CNL/CML/ICL), PDI2 reports an
    invalid count, force the correct hardware-supported value

    This may have to be revisited with platform-specific values if the
    hardware changes, but for now this is good enough.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20190806005522.22642-3-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5829ee5a5c98ea988dfb0c1ee94982989a24fe8f
Author: Will Deacon <will@kernel.org>
Date:   Thu Aug 8 16:51:00 2019 +0100

    ARM: 8898/1: mm: Don't treat faults reported from cache maintenance as writes

    [ Upstream commit 834020366da9ab3fb87d1eb9a3160eb22dbed63a ]

    Translation faults arising from cache maintenance instructions are
    rather unhelpfully reported with an FSR value where the WnR field is set
    to 1, indicating that the faulting access was a write. Since cache
    maintenance instructions on 32-bit ARM do not require any particular
    permissions, this can cause our private 'cacheflush' system call to fail
    spuriously if a translation fault is generated due to page aging when
    targetting a read-only VMA.

    In this situation, we will return -EFAULT to userspace, although this is
    unfortunately suppressed by the popular '__builtin___clear_cache()'
    intrinsic provided by GCC, which returns void.

    Although it's tempting to write this off as a userspace issue, we can
    actually do a little bit better on CPUs that support LPAE, even if the
    short-descriptor format is in use. On these CPUs, cache maintenance
    faults additionally set the CM field in the FSR, which we can use to
    suppress the write permission checks in the page fault handler and
    succeed in performing cache maintenance to read-only areas even in the
    presence of a translation fault.

    Reported-by: Orion Hodson <oth@google.com>
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6b0262aeaaab8cd58ce9d6c1fc3761e2e73700d8
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Thu Jun 13 15:43:20 2019 +0200

    mips/atomic: Fix smp_mb__{before,after}_atomic()

    [ Upstream commit 42344113ba7a1ed7b5654cd5270af0d5698d8521 ]

    Recent probing at the Linux Kernel Memory Model uncovered a
    'surprise'. Strongly ordered architectures where the atomic RmW
    primitive implies full memory ordering and
    smp_mb__{before,after}_atomic() are a simple barrier() (such as MIPS
    without WEAK_REORDERING_BEYOND_LLSC) fail for:

            *x = 1;
            atomic_inc(u);
            smp_mb__after_atomic();
            r0 = *y;

    Because, while the atomic_inc() implies memory order, it
```

```
    (surprisingly) does not provide a compiler barrier. This then allows
    the compiler to re-order like so:

            atomic_inc(u);
            *x = 1;
            smp_mb__after_atomic();
            r0 = *y;

    Which the CPU is then allowed to re-order (under TSO rules) like:

            atomic_inc(u);
            r0 = *y;
            *x = 1;

    And this very much was not intended. Therefore strengthen the atomic
    RmW ops to include a compiler barrier.

    Reported-by: Andrea Parri <andrea.parri@amarulasolutions.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Paul Burton <paul.burton@mips.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 528e1176fa7d377102c5e2f58baee84e74563fee
Author: Miroslav Benes <mbenes@suse.cz>
Date:   Fri Jul 19 14:28:39 2019 +0200

    livepatch: Nullify obj->mod in klp_module_coming()'s error path

    [ Upstream commit 4ff96fb52c6964ad42e0a878be8f86a2e8052ddd ]

    klp_module_coming() is called for every module appearing in the system.
    It sets obj->mod to a patched module for klp_object obj. Unfortunately
    it leaves it set even if an error happens later in the function and the
    patched module is not allowed to be loaded.

    klp_is_object_loaded() uses obj->mod variable and could currently give a
    wrong return value. The bug is probably harmless as of now.

    Signed-off-by: Miroslav Benes <mbenes@suse.cz>
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Acked-by: Josh Poimboeuf <jpoimboe@redhat.com>
    Signed-off-by: Petr Mladek <pmladek@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 318c2962506dc696781bb017f034cc75b336587c
Author: Xiaowei Bao <xiaowei.bao@nxp.com>
Date:   Wed Aug 14 10:03:29 2019 +0800

    PCI: layerscape: Add the bar_fixed_64bit property to the endpoint driver

    [ Upstream commit fd5d16531a39322c3d7433d9f8a36203c9aaeddc ]

    The layerscape PCIe controller have 4 BARs.

     BAR0 and BAR1 are 32bit, BAR2 and BAR4 are 64bit and that's a
    fixed hardware configuration.

    Set the bar_fixed_64bit variable accordingly.

    Signed-off-by: Xiaowei Bao <xiaowei.bao@nxp.com>
    [lorenzo.pieralisi@arm.com: commit log]
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Acked-by: Kishon Vijay Abraham I <kishon@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 69e6ea69d18a5def002ff752157fc38ccaf0d677
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Fri Jul 12 08:53:19 2019 -0700

    PCI: pci-hyperv: Fix build errors on non-SYSFS config

    [ Upstream commit f58ba5e3f6863ea4486952698898848a6db726c2 ]

    Fix build errors when building almost-allmodconfig but with SYSFS
    not set (not enabled). Fixes these build errors:

    ERROR: "pci_destroy_slot" [drivers/pci/controller/pci-hyperv.ko] undefined!
    ERROR: "pci_create_slot" [drivers/pci/controller/pci-hyperv.ko] undefined!

    drivers/pci/slot.o is only built when SYSFS is enabled, so
    pci-hyperv.o has an implicit dependency on SYSFS.
    Make that explicit.

    Also, depending on X86 && X86_64 is not needed, so just change that
    to depend on X86_64.

    Fixes: a15f2c08c708 ("PCI: hv: support reporting serial number as slot information")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Haiyang Zhang <haiyangz@microsoft.com>
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: Jake Oshins <jakeo@microsoft.com>
    Cc: "K. Y. Srinivasan" <kys@microsoft.com>
    Cc: Haiyang Zhang <haiyangz@microsoft.com>
    Cc: Stephen Hemminger <sthemmin@microsoft.com>
    Cc: Stephen Hemminger <stephen@networkplumber.org>
    Cc: Sasha Levin <sashal@kernel.org>
    Cc: Bjorn Helgaas <bhelgaas@google.com>
    Cc: linux-pci@vger.kernel.org
    Cc: linux-hyperv@vger.kernel.org
    Cc: Dexuan Cui <decui@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0a25ab4e0a2bcbc06e8a5f2e715ec84808d1f60
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Thu Jun 13 15:43:19 2019 +0200

    mips/atomic: Fix loongson_llsc_mb() wreckage

    [ Upstream commit 1c6c1ca318585f1096d4d04bc722297c85e9fb8a ]

    The comment describing the loongson_llsc_mb() reorder case doesn't
    make any sense what so ever. Instruction re-ordering is not an SMP
    artifact, but rather a CPU local phenomenon. Clarify the comment by
    explaining that these issue cause a coherence fail.

    For the branch speculation case; if futex_atomic_cmpxchg_inatomic()
    needs one at the bne branch target, then surely the normal
    __cmpxch_asm() implementation does too. We cannot rely on the
    barriers from cmpxchg() because cmpxchg_local() is implemented with
    the same macro, and branch prediction and speculation are, too, CPU
    local.

    Fixes: e02e07e3127d ("MIPS: Loongson: Introduce and use loongson_llsc_mb()")
    Cc: Huacai Chen <chenhc@lemote.com>
    Cc: Huang Pei <huangpei@loongson.cn>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Paul Burton <paul.burton@mips.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8d46fca32c2a3b6022aba6d6d93ace1c18d31f68
Author: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
Date:   Mon Aug 12 09:36:11 2019 +0300

    rtc: bd70528: fix driver dependencies

    [ Upstream commit 41a8e19f47dfe9154c56b35842700ad38a7c33e0 ]

    With CONFIG_BD70528_WATCHDOG=m, a built-in rtc driver cannot call
    into the low-level functions that are part of the watchdog module:

    drivers/rtc/rtc-bd70528.o: In function `bd70528_set_time':
    rtc-bd70528.c:(.text+0x22c): undefined reference to `bd70528_wdt_lock'
    rtc-bd70528.c:(.text+0x2a8): undefined reference to `bd70528_wdt_unlock'
    drivers/rtc/rtc-bd70528.o: In function `bd70528_set_rtc_based_timers':
    rtc-bd70528.c:(.text+0x50c): undefined reference to `bd70528_wdt_set'

    Add a Kconfig dependency which forces RTC to be a module if watchdog is a
    module. If watchdog is not compiled at all the stub functions for watchdog
    control are used. compiling the RTC without watchdog is fine.

    Fixes: 32a4a4ebf768 ("rtc: bd70528: Initial support for ROHM bd70528 RTC")
    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Reported-by: kbuild test robot <lkp@intel.com>
    Signed-off-by: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
    Link:
https://lore.kernel.org/r/84462e01e43d39024948a3bdd24087ff87dc2255.1565591387.git.matti.vait
tinen@fi.rohmeurope.com
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit afa6eca42495e3f5fecaf771a05cdb9f56f439e8
Author: Jason Gerecke <killertofu@gmail.com>
Date:   Fri Aug 16 11:54:26 2019 -0700

    HID: wacom: Fix several minor compiler warnings

    [ Upstream commit 073b50bccbbf99a3b79a1913604c656d0e1a56c9 ]

    Addresses a few issues that were noticed when compiling with non-default
    warnings enabled. The trimmed-down warnings in the order they are fixed
    below are:

    * declaration of 'size' shadows a parameter

    * '%s' directive output may be truncated writing up to 5 bytes into a
      region of size between 1 and 64

    * pointer targets in initialization of 'char *' from 'unsigned char *'
      differ in signedness

    * left shift of negative value

    Signed-off-by: Jason Gerecke <jason.gerecke@wacom.com>
    Reviewed-by: Aaron Armstrong Skomra <aaron.skomra@wacom.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 09af73dac76dbc82548ec247954a20619ff1b2df
Author: Nishka Dasgupta <nishkadg.linux@gmail.com>
Date:   Wed Jul 24 13:54:12 2019 +0530

    PCI: tegra: Fix OF node reference leak

    [ Upstream commit 9e38e690ace3e7a22a81fc02652fc101efb340cf ]
```

```
    Each iteration of for_each_child_of_node() executes of_node_put() on the
    previous node, but in some return paths in the middle of the loop
    of_node_put() is missing thus causing a reference leak.

    Hence stash these mid-loop return values in a variable 'err' and add a
    new label err_node_put which executes of_node_put() on the previous node
    and returns 'err' on failure.

    Change mid-loop return statements to point to jump to this label to
    fix the reference leak.

    Issue found with Coccinelle.

    Signed-off-by: Nishka Dasgupta <nishkadg.linux@gmail.com>
    [lorenzo.pieralisi@arm.com: rewrote commit log]
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6be42301c577bb24ff1091d474803af8b26ab31b
Author: Kai-Heng Feng <kai.heng.feng@canonical.com>
Date:   Fri Jul 5 12:55:03 2019 +0800

    mfd: intel-lpss: Remove D3cold delay

    [ Upstream commit 76380a607ba0b28627c9b4b55cd47a079a59624b ]

    Goodix touchpad may drop its first couple input events when
    i2c-designware-platdrv and intel-lpss it connects to took too long to
    runtime resume from runtime suspended state.

    This issue happens becuase the touchpad has a rather small buffer to
    store up to 13 input events, so if the host doesn't read those events in
    time (i.e. runtime resume takes too long), events are dropped from the
    touchpad's buffer.

    The bottleneck is D3cold delay it waits when transitioning from D3cold
    to D0, hence remove the delay to make the resume faster. I've tested
    some systems with intel-lpss and haven't seen any regression.

    Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=202683
    Signed-off-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22a5fd6c1a0a0a3c3f619780ec43461a331e18f2
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Tue Aug 13 12:03:01 2019 +0200

    i2c-cht-wc: Fix lockdep warning

    [ Upstream commit 232219b9a464c2479c98aa589acb1bd3383ae9d6 ]

    When the kernel is build with lockdep support and the i2c-cht-wc driver is
    used, the following warning is shown:

    [   66.674334] ============================================================
    [   66.674337] WARNING: possible circular locking dependency detected
    [   66.674340] 5.3.0-rc4+ #83 Not tainted
    [   66.674342] ------------------------------------------------------
    [   66.674345] systemd-udevd/1232 is trying to acquire lock:
    [   66.674349] 00000000a74dab07 (intel_soc_pmic_chtwc:167:(&cht_wc_regmap_cfg)->lock)
{+.+.}, at: regmap_write+0x31/0x70
    [   66.674360]
```

```
                    but task is already holding lock:
     [   66.674362] 00000000d44a85b7 (i2c_register_adapter){+.+.}, at:
  i2c_smbus_xfer+0x49/0xf0
     [   66.674370]
                    which lock already depends on the new lock.

     [   66.674371]
                    the existing dependency chain (in reverse order) is:
     [   66.674374]
                    -> #1 (i2c_register_adapter){+.+.}:
     [   66.674381]        rt_mutex_lock_nested+0x46/0x60
     [   66.674384]        i2c_smbus_xfer+0x49/0xf0
     [   66.674387]        i2c_smbus_read_byte_data+0x45/0x70
     [   66.674391]        cht_wc_byte_reg_read+0x35/0x50
     [   66.674394]        _regmap_read+0x63/0x1a0
     [   66.674396]        _regmap_update_bits+0xa8/0xe0
     [   66.674399]        regmap_update_bits_base+0x63/0xa0
     [   66.674403]        regmap_irq_update_bits.isra.0+0x3b/0x50
     [   66.674406]        regmap_add_irq_chip+0x592/0x7a0
     [   66.674409]        devm_regmap_add_irq_chip+0x89/0xed
     [   66.674412]        cht_wc_probe+0x102/0x158
     [   66.674415]        i2c_device_probe+0x95/0x250
     [   66.674419]        really_probe+0xf3/0x380
     [   66.674422]        driver_probe_device+0x59/0xd0
     [   66.674425]        device_driver_attach+0x53/0x60
     [   66.674428]        __driver_attach+0x92/0x150
     [   66.674431]        bus_for_each_dev+0x7d/0xc0
     [   66.674434]        bus_add_driver+0x14d/0x1f0
     [   66.674437]        driver_register+0x6d/0xb0
     [   66.674440]        i2c_register_driver+0x45/0x80
     [   66.674445]        do_one_initcall+0x60/0x2f4
     [   66.674450]        kernel_init_freeable+0x20d/0x2b4
     [   66.674453]        kernel_init+0xa/0x10c
     [   66.674457]        ret_from_fork+0x3a/0x50
     [   66.674459]
                    -> #0 (intel_soc_pmic_chtwc:167:(&cht_wc_regmap_cfg)->lock){+.+.}:
     [   66.674465]        __lock_acquire+0xe07/0x1930
     [   66.674468]        lock_acquire+0x9d/0x1a0
     [   66.674472]        __mutex_lock+0xa8/0x9a0
     [   66.674474]        regmap_write+0x31/0x70
     [   66.674480]        cht_wc_i2c_adap_smbus_xfer+0x72/0x240 [i2c_cht_wc]
     [   66.674483]        __i2c_smbus_xfer+0x1a3/0x640
     [   66.674486]        i2c_smbus_xfer+0x67/0xf0
     [   66.674489]        i2c_smbus_read_byte_data+0x45/0x70
     [   66.674494]        bq24190_probe+0x26b/0x410 [bq24190_charger]
     [   66.674497]        i2c_device_probe+0x189/0x250
     [   66.674500]        really_probe+0xf3/0x380
     [   66.674503]        driver_probe_device+0x59/0xd0
     [   66.674506]        device_driver_attach+0x53/0x60
     [   66.674509]        __driver_attach+0x92/0x150
     [   66.674512]        bus_for_each_dev+0x7d/0xc0
     [   66.674515]        bus_add_driver+0x14d/0x1f0
     [   66.674518]        driver_register+0x6d/0xb0
     [   66.674521]        i2c_register_driver+0x45/0x80
     [   66.674524]        do_one_initcall+0x60/0x2f4
     [   66.674528]        do_init_module+0x5c/0x230
     [   66.674531]        load_module+0x2707/0x2a20
     [   66.674534]        __do_sys_init_module+0x188/0x1b0
     [   66.674537]        do_syscall_64+0x5c/0xb0
     [   66.674541]        entry_SYSCALL_64_after_hwframe+0x49/0xbe
     [   66.674543]
                    other info that might help us debug this:

     [   66.674545]  Possible unsafe locking scenario:
```

```
     [   66.674547]        CPU0                    CPU1
     [   66.674548]        ----                    ----
     [   66.674550]   lock(i2c_register_adapter);
     [   66.674553]                                lock(intel_soc_pmic_chtwc:167:
  (&cht_wc_regmap_cfg)->lock);
     [   66.674556]                                lock(i2c_register_adapter);
     [   66.674559]   lock(intel_soc_pmic_chtwc:167:(&cht_wc_regmap_cfg)->lock);
     [   66.674561]
                    *** DEADLOCK ***
```

    The problem is that the CHT Whiskey Cove PMIC's builtin i2c-adapter is
    itself a part of an i2c-client (the PMIC). This means that transfers done
    through it take adapter->bus_lock twice, once for the parent i2c-adapter
    and once for its own bus_lock. Lockdep does not like this nested locking.

    To make lockdep happy in the case of busses with muxes, the i2c-core's
    i2c_adapter_lock_bus function calls:

     rt_mutex_lock_nested(&adapter->bus_lock, i2c_adapter_depth(adapter));

    But i2c_adapter_depth only works when the direct parent of the adapter is
    another adapter, as it is only meant for muxes. In this case there is an
    i2c-client and MFD instantiated platform_device in the parent->child chain
    between the 2 devices.

    This commit overrides the default i2c_lock_operations, passing a hardcoded
    depth of 1 to rt_mutex_lock_nested, making lockdep happy.

    Note that if there were to be a mux attached to the i2c-wc-cht adapter,
    this would break things again since the i2c-mux code expects the
    root-adapter to have a locking depth of 0. But the i2c-wc-cht adapter
    always has only 1 client directly attached in the form of the charger IC
    paired with the CHT Whiskey Cove PMIC.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Wolfram Sang <wsa@the-dreams.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8803f0ad3ba1ba7f0d9f56c594f8bc77d6ee1f99
Author: Nathan Chancellor <natechancellor@gmail.com>
Date: Sun Aug 11 20:31:20 2019 -0700

    MIPS: tlbex: Explicitly cast _PAGE_NO_EXEC to a boolean

    [Upstream commit c59ae0a1055127dd3828a88e111a0db59b254104]

    clang warns:

```
    arch/mips/mm/tlbex.c:634:19: error: use of logical '&&' with constant
    operand [-Werror,-Wconstant-logical-operand]
            if (cpu_has_rixi && _PAGE_NO_EXEC) {
                             ^  ~~~~~~~~~~~~~~
    arch/mips/mm/tlbex.c:634:19: note: use '&' for a bitwise operation
            if (cpu_has_rixi && _PAGE_NO_EXEC) {
                             ^~
                             &
    arch/mips/mm/tlbex.c:634:19: note: remove constant to silence this
    warning
            if (cpu_has_rixi && _PAGE_NO_EXEC) {
                            ~^~~~~~~~~~~~~~~~
    1 error generated.
```

    Explicitly cast this value to a boolean so that clang understands we
    intend for this to be a non-zero value.

```
    Fixes: 00bf1c691d08 ("MIPS: tlbex: Avoid placing software PTE bits in Entry* PFN
fields")
    Link: https://github.com/ClangBuiltLinux/linux/issues/609
    Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
    Signed-off-by: Paul Burton <paul.burton@mips.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: linux-mips@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Cc: clang-built-linux@googlegroups.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff0cdd0d33d94b18e59c0501f7c75081cf8e6a71
Author: Nathan Chancellor <natechancellor@gmail.com>
Date:   Sun Aug 11 20:31:16 2019 -0700

    MIPS: Don't use bc_false uninitialized in __mm_isBranchInstr

    [ Upstream commit c2869aafe7191d366d74c55cb8a93c6d0baba317 ]

    clang warns:

    arch/mips/kernel/branch.c:148:8: error: variable 'bc_false' is used
    uninitialized whenever switch case is taken
    [-Werror,-Wsometimes-uninitialized]
                    case mm_bc2t_op:
                         ^~~~~~~~~~
    arch/mips/kernel/branch.c:157:8: note: uninitialized use occurs here
                            if (bc_false)
                                ^~~~~~~~
    arch/mips/kernel/branch.c:149:8: error: variable 'bc_false' is used
    uninitialized whenever switch case is taken
    [-Werror,-Wsometimes-uninitialized]
                    case mm_bc1t_op:
                         ^~~~~~~~~~
    arch/mips/kernel/branch.c:157:8: note: uninitialized use occurs here
                            if (bc_false)
                                ^~~~~~~~
    arch/mips/kernel/branch.c:142:4: note: variable 'bc_false' is declared
    here
                            int bc_false = 0;
                            ^
    2 errors generated.

    When mm_bc1t_op and mm_bc2t_op are taken, the bc_false initialization
    does not happen, which leads to a garbage value upon use, as illustrated
    below with a small sample program.

    $ mipsel-linux-gnu-gcc --version | head -n1
    mipsel-linux-gnu-gcc (Debian 8.3.0-2) 8.3.0

    $ clang --version | head -n1
    ClangBuiltLinux clang version 9.0.0 (git://github.com/llvm/llvm-project
    544315b4197034a3be8acd12cba56a75fb1f08dc) (based on LLVM 9.0.0svn)

    $ cat test.c
     #include <stdio.h>

     static void switch_scoped(int opcode)
     {
             switch (opcode) {
             case 1:
             case 2: {
                     int bc_false = 0;
```

```
                bc_false = 4;
        case 3:
        case 4:
                printf("\t* switch scoped bc_false = %d\n", bc_false);
        }
        }
}

static void function_scoped(int opcode)
{
        int bc_false = 0;

        switch (opcode) {
        case 1:
        case 2: {
                bc_false = 4;
        case 3:
        case 4:
                printf("\t* function scoped bc_false = %d\n", bc_false);
        }
        }
}

int main(void)
{
        int opcode;

        for (opcode = 1; opcode < 5; opcode++) {
                printf("opcode = %d:\n", opcode);
                switch_scoped(opcode);
                function_scoped(opcode);
                printf("\n");
        }

        return 0;
}

$ mipsel-linux-gnu-gcc -std=gnu89 -static test.c && \
  qemu-mipsel a.out
opcode = 1:
        * switch scoped bc_false = 4
        * function scoped bc_false = 4

opcode = 2:
        * switch scoped bc_false = 4
        * function scoped bc_false = 4

opcode = 3:
        * switch scoped bc_false = 2147483004
        * function scoped bc_false = 0

opcode = 4:
        * switch scoped bc_false = 2147483004
        * function scoped bc_false = 0

$ clang -std=gnu89 --target=mipsel-linux-gnu -m32 -static test.c && \
  qemu-mipsel a.out
opcode = 1:
        * switch scoped bc_false = 4
        * function scoped bc_false = 4

opcode = 2:
        * switch scoped bc_false = 4
        * function scoped bc_false = 4
```

```
    opcode = 3:
            * switch scoped bc_false = 2147483004
            * function scoped bc_false = 0

    opcode = 4:
            * switch scoped bc_false = 2147483004
            * function scoped bc_false = 0

    Move the definition up so that we get the right behavior and mark it
    __maybe_unused as it will not be used when CONFIG_MIPS_FP_SUPPORT
    isn't enabled.

    Fixes: 6a1cc218b9cc ("MIPS: branch: Remove FP branch handling when
CONFIG_MIPS_FP_SUPPORT=n")
    Link: https://github.com/ClangBuiltLinux/linux/issues/603
    Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
    Signed-off-by: Paul Burton <paul.burton@mips.com>
    Cc: Ralf Baechle <ralf@linux-mips.org>
    Cc: James Hogan <jhogan@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: linux-mips@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Cc: clang-built-linux@googlegroups.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9d4d767ff325d919d554295b4ef2f4e8dd7e787
Author: Zhou Yanjie <zhouyanjie@zoho.com>
Date:   Fri Aug 2 16:27:37 2019 +0800

    MIPS: Ingenic: Disable broken BTB lookup optimization.

    [ Upstream commit 053951dda71ecb4b554a2cdbe26f5f6f9bee9dd2 ]

    In order to further reduce power consumption, the XBurst core
    by default attempts to avoid branch target buffer lookups by
    detecting & special casing loops. This feature will cause
    BogoMIPS and lpj calculate in error. Set cp0 config7 bit 4 to
    disable this feature.

    Signed-off-by: Zhou Yanjie <zhouyanjie@zoho.com>
    Signed-off-by: Paul Burton <paul.burton@mips.com>
    Cc: linux-mips@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Cc: ralf@linux-mips.org
    Cc: paul@crapouillou.net
    Cc: jhogan@kernel.org
    Cc: malat@debian.org
    Cc: gregkh@linuxfoundation.org
    Cc: tglx@linutronix.de
    Cc: allison@lohutok.net
    Cc: syq@debian.org
    Cc: chenhc@lemote.com
    Cc: jiaxun.yang@flygoat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9fd765009687226caa8de88dd8a963efd3245f64
Author: zhangyi (F) <yi.zhang@huawei.com>
Date:   Wed Aug 28 11:13:24 2019 -0400

    ext4: fix potential use after free after remounting with noblock_validity

    [ Upstream commit 7727ae52975d4f4ef7ff69ed8e6e25f6a4168158 ]

    Remount process will release system zone which was allocated before if
```

```
    "noblock_validity" is specified. If we mount an ext4 file system to two
    mountpoints with default mount options, and then remount one of them
    with "noblock_validity", it may trigger a use after free problem when
    someone accessing the other one.

     # mount /dev/sda foo
     # mount /dev/sda bar

    User access mountpoint "foo"   |   Remount mountpoint "bar"
                                   |
    ext4_map_blocks()              |   ext4_remount()
    check_block_validity()         |   ext4_setup_system_zone()
    ext4_data_block_valid()        |   ext4_release_system_zone()
                                   |   free system_blks rb nodes
    access system_blks rb nodes    |
    trigger use after free         |

    This problem can also be reproduced by one mountpint, At the same time,
    add_system_zone() can get called during remount as well so there can be
    racing ext4_data_block_valid() reading the rbtree at the same time.

    This patch add RCU to protect system zone from releasing or building
    when doing a remount which inverse current "noblock_validity" mount
    option. It assign the rbtree after the whole tree was complete and
    do actual freeing after rcu grace period, avoid any intermediate state.

    Reported-by: syzbot+1e470567330b7ad711d5@syzkaller.appspotmail.com
    Signed-off-by: zhangyi (F) <yi.zhang@huawei.com>
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Reviewed-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 893f29bad19a49e8b72504faf6f621577c6e8ab7
Author: Chao Yu <yuchao0@huawei.com>
Date:   Thu Jul 18 16:39:59 2019 +0800

    f2fs: fix to drop meta/node pages during umount

    [ Upstream commit a8933b6b68f775b5774e7b075447fae13f4d01fe ]

    As reported in bugzilla:

    https://bugzilla.kernel.org/show_bug.cgi?id=204193

    A null pointer dereference bug is triggered in f2fs under kernel-5.1.3.

     kasan_report.cold+0x5/0x32
     f2fs_write_end_io+0x215/0x650
     bio_endio+0x26e/0x320
     blk_update_request+0x209/0x5d0
     blk_mq_end_request+0x2e/0x230
     lo_complete_rq+0x12c/0x190
     blk_done_softirq+0x14a/0x1a0
     __do_softirq+0x119/0x3e5
     irq_exit+0x94/0xe0
     call_function_single_interrupt+0xf/0x20

    During umount, we will access NULL sbi->node_inode pointer in
    f2fs_write_end_io():

            f2fs_bug_on(sbi, page->mapping == NODE_MAPPING(sbi) &&
                                    page->index != nid_of_node(page));

    The reason is if disable_checkpoint mount option is on, meta dirty
    pages can remain during umount, and then be flushed by iput() of
```

```
    meta_inode, however node_inode has been iput()ed before
    meta_inode's iput().

    Since checkpoint is disabled, all meta/node datas are useless and
    should be dropped in next mount, so in umount, let's adjust
    drop_inode() to give a hint to iput_final() to drop all those dirty
    datas correctly.

    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5d9658aa7b0bbf2ac8a74955e4846de1a9fd6246
Author: Chris Wilson <chris@chris-wilson.co.uk>
Date:   Mon Aug 12 16:42:47 2019 +0100

    dma-buf/sw_sync: Synchronize signal vs syncpt free

    [ Upstream commit d3c6dd1fb30d3853c2012549affe75c930f4a2f9 ]

    During release of the syncpt, we remove it from the list of syncpt and
    the tree, but only if it is not already been removed. However, during
    signaling, we first remove the syncpt from the list. So, if we
    concurrently free and signal the syncpt, the free may decide that it is
    not part of the tree and immediately free itself -- meanwhile the
    signaler goes on to use the now freed datastructure.

    In particular, we get struck by commit 0e2f733addbf ("dma-buf: make
    dma_fence structure a bit smaller v2") as the cb_list is immediately
    clobbered by the kfree_rcu.

    v2: Avoid calling into timeline_fence_release() from under the spinlock

    Bugzilla: https://bugs.freedesktop.org/show_bug.cgi?id=111381
    Fixes: d3862e44daa7 ("dma-buf/sw-sync: Fix locking around sync_timeline lists")
    References: 0e2f733addbf ("dma-buf: make dma_fence structure a bit smaller v2")
    Signed-off-by: Chris Wilson <chris@chris-wilson.co.uk>
    Cc: Sumit Semwal <sumit.semwal@linaro.org>
    Cc: Sean Paul <seanpaul@chromium.org>
    Cc: Gustavo Padovan <gustavo@padovan.org>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: <stable@vger.kernel.org> # v4.14+
    Acked-by: Christian König <christian.koenig@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190812154247.20508-1-chris@chris-
wilson.co.uk
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01749cb7230132ea4aaf0f129ced0f364059795a
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Thu Aug 1 15:38:14 2019 -0700

    scsi: core: Reduce memory required for SCSI logging

    [ Upstream commit dccc96abfb21dc19d69e707c38c8ba439bba7160 ]

    The data structure used for log messages is so large that it can cause a
    boot failure. Since allocations from that data structure can fail anyway,
    use kmalloc() / kfree() instead of that data structure.

    See also https://bugzilla.kernel.org/show_bug.cgi?id=204119.
    See also commit ded85c193a39 ("scsi: Implement per-cpu logging buffer") # v4.0.

    Reported-by: Jan Palus <jpalus@fastmail.com>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Hannes Reinecke <hare@suse.com>
```

```
    Cc: Johannes Thumshirn <jthumshirn@suse.de>
    Cc: Ming Lei <ming.lei@redhat.com>
    Cc: Jan Palus <jpalus@fastmail.com>
    Signed-off-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 97e9c0a70df27b979968f802aeaa7d2088ceac86
Author: Chunyan Zhang <chunyan.zhang@unisoc.com>
Date:   Thu Sep 5 18:30:09 2019 +0800

    clk: sprd: add missing kfree

    [ Upstream commit 5e75ea9c67433a065b0e8595ad3c91c7c0ca0d2d ]

    The number of config registers for different pll clocks probably are not
    same, so we have to use malloc, and should free the memory before return.

    Fixes: 3e37b005580b ("clk: sprd: add adjustable pll support")
    Signed-off-by: Chunyan Zhang <chunyan.zhang@unisoc.com>
    Signed-off-by: Chunyan Zhang <zhang.lyra@gmail.com>
    Link: https://lkml.kernel.org/r/20190905103009.27166-1-zhang.lyra@gmail.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fc4d0a10863463c6f3734e6d50d432872168832b
Author: Jorge Ramirez-Ortiz <jorge.ramirez-ortiz@linaro.org>
Date:   Thu Aug 29 10:27:58 2019 +0200

    mbox: qcom: add APCS child device for QCS404

    [ Upstream commit 78c86458a440ff356073c21b568cb58ddb67b82b ]

    There is clock controller functionality in the APCS hardware block of
    qcs404 devices similar to msm8916.

    Co-developed-by: Niklas Cassel <niklas.cassel@linaro.org>
    Signed-off-by: Niklas Cassel <niklas.cassel@linaro.org>
    Signed-off-by: Jorge Ramirez-Ortiz <jorge.ramirez-ortiz@linaro.org>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Reviewed-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5af20ee3e3bcd58785ca3c0ec0d5038030b65673
Author: Ganesh Goudar <ganeshgr@linux.ibm.com>
Date:   Wed Sep 4 13:29:49 2019 +0530

    powerpc: dump kernel log before carrying out fadump or kdump

    [ Upstream commit e7ca44ed3ba77fc26cf32650bb71584896662474 ]

    Since commit 4388c9b3a6ee ("powerpc: Do not send system reset request
    through the oops path"), pstore dmesg file is not updated when dump is
    triggered from HMC. This commit modified system reset (sreset) handler
    to invoke fadump or kdump (if configured), without pushing dmesg to
    pstore. This leaves pstore to have old dmesg data which won't be much
    of a help if kdump fails to capture the dump. This patch fixes that by
    calling kmsg_dump() before heading to fadump ot kdump.

    Fixes: 4388c9b3a6ee ("powerpc: Do not send system reset request through the oops path")
    Reviewed-by: Mahesh Salgaonkar <mahesh@linux.vnet.ibm.com>
    Reviewed-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Ganesh Goudar <ganeshgr@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
```

```
    Link: https://lore.kernel.org/r/20190904075949.15607-1-ganeshgr@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 21c37928be9b7b87e4c4483ef9fcd09050367be4
Author: Bjorn Andersson <bjorn.andersson@linaro.org>
Date:   Thu Sep 12 19:40:29 2019 -0700

    clk: Make clk_bulk_get_all() return a valid "id"

    [ Upstream commit 7f81c2426587b34bf73e643c1a6d080dfa14cf8a ]

    The adreno driver expects the "id" field of the returned clk_bulk_data
    to be filled in with strings from the clock-names property.

    But due to the use of kmalloc_array() in of_clk_bulk_get_all() it
    receives a list of bogus pointers instead.

    Zero-initialize the "id" field and attempt to populate with strings from
    the clock-names property to resolve both these issues.

    Fixes: 616e45df7c4a ("clk: add new APIs to operate on all available clocks")
    Fixes: 8e3e791d20d2 ("drm/msm: Use generic bulk clock function")
    Cc: Dong Aisheng <aisheng.dong@nxp.com>
    Cc: Jordan Crouse <jcrouse@codeaurora.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lkml.kernel.org/r/20190913024029.2640-1-bjorn.andersson@linaro.org
    Reviewed-by: Jordan Crouse <jcrouse@codeaurora.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97322b0539716975486a1d6c20aa99c9b205a94b
Author: Peng Fan <peng.fan@nxp.com>
Date:   Mon Sep 9 03:39:39 2019 +0000

    clk: imx: clk-pll14xx: unbypass PLL by default

    [ Upstream commit a9aa8306074d9519dd6e5fdf07240b01bac72e04 ]

    When registering the PLL, unbypass the PLL.
    The PLL has two bypass control bit, BYPASS and EXT_BYPASS.
    we will expose EXT_BYPASS to clk driver for mux usage, and keep
    BYPASS inside pll14xx usage. The PLL has a restriction that
    when M/P change, need to RESET/BYPASS pll to avoid glitch, so
    we could not expose BYPASS.

    To make it easy for clk driver usage, unbypass PLL which does
    not hurt current function.

    Fixes: 8646d4dcc7fb ("clk: imx: Add PLLs driver for imx8mm soc")
    Reviewed-by: Leonard Crestez <leonard.crestez@nxp.com>
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Link: https://lkml.kernel.org/r/1568043491-20680-3-git-send-email-peng.fan@nxp.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 292a0dc34cd869314f0169511b2f0f82c4c8b7ad
Author: Peng Fan <peng.fan@nxp.com>
Date:   Mon Sep 9 03:39:34 2019 +0000

    clk: imx: pll14xx: avoid glitch when set rate

    [ Upstream commit dee1bc9c23cd41fe32549c0adbe6cb57cab02282 ]

    According to PLL1443XA and PLL1416X spec,
    "When BYPASS is 0 and RESETB is changed from 0 to 1, FOUT starts to
```

```
    output unstable clock until lock time passes. PLL1416X/PLL1443XA may
    generate a glitch at FOUT."

    So set BYPASS when RESETB is changed from 0 to 1 to avoid glitch.
    In the end of set rate, BYPASS will be cleared.

    When prepare clock, also need to take care to avoid glitch. So
    we also follow Spec to set BYPASS before RESETB changed from 0 to 1.
    And add a check if the RESETB is already 0, directly return 0;

    Fixes: 8646d4dcc7fb ("clk: imx: Add PLLs driver for imx8mm soc")
    Reviewed-by: Leonard Crestez <leonard.crestez@nxp.com>
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Link: https://lkml.kernel.org/r/1568043491-20680-2-git-send-email-peng.fan@nxp.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b15a76dde4c7b4cd7c4c14cd1ff66e569491603d
Author: Eugen Hristev <eugen.hristev@microchip.com>
Date:   Mon Sep 9 15:30:34 2019 +0000

    clk: at91: select parent if main oscillator or bypass is enabled

    [ Upstream commit 69a6bcde7fd3fe6f3268ce26f31d9d9378384c98 ]

    Selecting the right parent for the main clock is done using only
    main oscillator enabled bit.
    In case we have this oscillator bypassed by an external signal (no driving
    on the XOUT line), we still use external clock, but with BYPASS bit set.
    So, in this case we must select the same parent as before.
    Create a macro that will select the right parent considering both bits from
    the MOR register.
    Use this macro when looking for the right parent.

    Signed-off-by: Eugen Hristev <eugen.hristev@microchip.com>
    Link: https://lkml.kernel.org/r/1568042692-11784-2-git-send-email-
eugen.hristev@microchip.com
    Acked-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Reviewed-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1a5865508148a14f6275c37f0cbc33a180be62d
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Tue Sep 10 13:56:22 2019 +0200

    arm64: fix unreachable code issue with cmpxchg

    [ Upstream commit 920fdab7b3ce98c14c840261e364f490f3679a62 ]

    On arm64 build with clang, sometimes the __cmpxchg_mb is not inlined
    when CONFIG_OPTIMIZE_INLINING is set.
    Clang then fails a compile-time assertion, because it cannot tell at
    compile time what the size of the argument is:

    mm/memcontrol.o: In function `__cmpxchg_mb':
    memcontrol.c:(.text+0x1a4c): undefined reference to `__compiletime_assert_175'
    memcontrol.c:(.text+0x1a4c): relocation truncated to fit: R_AARCH64_CALL26 against
undefined symbol `__compiletime_assert_175'

    Mark all of the cmpxchg() style functions as __always_inline to
    ensure that the compiler can see the result.

    Acked-by: Nick Desaulniers <ndesaulniers@google.com>
    Reported-by: Nathan Chancellor <natechancellor@gmail.com>
```

```
    Link: https://github.com/ClangBuiltLinux/linux/issues/648
    Reviewed-by: Nathan Chancellor <natechancellor@gmail.com>
    Tested-by: Nathan Chancellor <natechancellor@gmail.com>
    Reviewed-by: Andrew Murray <andrew.murray@arm.com>
    Tested-by: Andrew Murray <andrew.murray@arm.com>
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b5582c3d5e129f2414b335a95fd9eefafcced86e
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun Sep 15 15:54:44 2019 +0200

    ARM: dts: dir685: Drop spi-cpol from the display

    [ Upstream commit 2a7326caab479ca257c4b9bd67db42d1d49079bf ]

    The D-Link DIR-685 had its clock polarity set as active
    low using the special SPI "spi-cpol" property.

    This is not correct: the datasheet clearly states:
    "Fix SCL to GND level when not in use" which is
    indicative that this line is active high.

    After a recent fix making the GPIO-based SPI driver
    force the clock line de-asserted at the beginning of
    each SPI transaction this reared its ugly head: now
    de-asserted was taken to mean the line should be
    driven high, but it should be driven low.

    Fix this up in the DTS file and the display works again.

    Link: https://lore.kernel.org/r/20190915135444.11066-1-linus.walleij@linaro.org
    Cc: Mark Brown <broonie@kernel.org>
    Fixes: 2922d1cc1696 ("spi: gpio: Add SPI_MASTER_GPIO_SS flag")
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd770fb86f998f15caeac5e9e9284a910db8db3b
Author: Bibby Hsieh <bibby.hsieh@mediatek.com>
Date:   Thu Aug 29 09:48:12 2019 +0800

    mailbox: mediatek: cmdq: clear the event in cmdq initial flow

    [ Upstream commit 6058f11870b8e6d4f5cc7b591097c00bf69a000d ]

    GCE hardware stored event information in own internal sysram,
    if the initial value in those sysram is not zero value
    it will cause a situation that gce can wait the event immediately
    after client ask gce to wait event but not really trigger the
    corresponding hardware.

    In order to make sure that the wait event function is
    exactly correct, we need to clear the sysram value in
    cmdq initial flow.

    Fixes: 623a6143a845 ("mailbox: mediatek: Add Mediatek CMDQ driver")

    Signed-off-by: Bibby Hsieh <bibby.hsieh@mediatek.com>
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit fcefc1b1761a8f40736e2739e03065c2b8ba2e23
Author: Otto Meier <gf435@gmx.net>
Date:   Tue Sep 10 18:01:41 2019 +0200

    pinctrl: meson-gxbb: Fix wrong pinning definition for uart_c

    [ Upstream commit cb0438e4436085d89706b5ccfce4d5da531253de ]

    Hi i tried to use the uart_C of the the odroid-c2.

    I enabled it in the dts file. During boot it crashed when the
    the sdcard slot is addressed.

    After long search in the net i found this:

    https://forum.odroid.com/viewtopic.php?f=139&t=25371&p=194370&hilit=uart_C#p177856

    After changing the pin definitions accordingly erverything works.
    Uart_c is functioning and sdcard ist working.

    Fixes: 6db0f3a8a04e46 ("pinctrl: amlogic: gxbb: add more UART pins")
    Signed-off-by: Otto Meier <gf435@gmx.net>
    Link: https://lore.kernel.org/r/1cc32a18-464d-5531-7a1c-084390e2ecb1@gmx.net
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22d3a69e0880584fbecf65b88d96e8184d23b7fa
Author: Nathan Lynch <nathanl@linux.ibm.com>
Date:   Tue Sep 10 17:52:44 2019 -0500

    powerpc/pseries: correctly track irq state in default idle

    [ Upstream commit 92c94dfb69e350471473fd3075c74bc68150879e ]

    prep_irq_for_idle() is intended to be called before entering
    H_CEDE (and it is used by the pseries cpuidle driver). However the
    default pseries idle routine does not call it, leading to mismanaged
    lazy irq state when the cpuidle driver isn't in use. Manifestations of
    this include:

    * Dropped IPIs in the time immediately after a cpu comes
      online (before it has installed the cpuidle handler), making the
      online operation block indefinitely waiting for the new cpu to
      respond.

    * Hitting this WARN_ON in arch_local_irq_restore():
            /*
             * We should already be hard disabled here. We had bugs
             * where that wasn't the case so let's dbl check it and
             * warn if we are wrong. Only do that when IRQ tracing
             * is enabled as mfmsr() can be costly.
             */
            if (WARN_ON_ONCE(mfmsr() & MSR_EE))
                    __hard_irq_disable();

    Call prep_irq_for_idle() from pseries_lpar_idle() and honor its
    result.

    Fixes: 363edbe2614a ("powerpc: Default arch idle could cede processor on pseries")
    Signed-off-by: Nathan Lynch <nathanl@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190910225244.25056-1-nathanl@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e04cfbae7d54e1d335e83e913cf442d542c229a6
```

```
Author: Stephen Boyd <swboyd@chromium.org>
Date:   Fri Aug 30 12:51:42 2019 -0700

    clk: qcom: gcc-sdm845: Use floor ops for sdcc clks

    [ Upstream commit 5e4b7e82d497580bc430576c4c9bce157dd72512 ]

    Some MMC cards fail to enumerate properly when inserted into an MMC slot
    on sdm845 devices. This is because the clk ops for qcom clks round the
    frequency up to the nearest rate instead of down to the nearest rate.
    For example, the MMC driver requests a frequency of 52MHz from
    clk_set_rate() but the qcom implementation for these clks rounds 52MHz
    up to the next supported frequency of 100MHz. The MMC driver could be
    modified to request clk rate ranges but for now we can fix this in the
    clk driver by changing the rounding policy for this clk to be round down
    instead of round up.

    Fixes: 06391eddb60a ("clk: qcom: Add Global Clock controller (GCC) driver for SDM845")
    Reported-by: Douglas Anderson <dianders@chromium.org>
    Cc: Taniya Das <tdas@codeaurora.org>
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Link: https://lkml.kernel.org/r/20190830195142.103564-1-swboyd@chromium.org
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5248e64491a91965ea0cb1ed8661abd07feb2d40
Author: Oliver O'Halloran <oohall@gmail.com>
Date:   Tue Sep 3 20:15:52 2019 +1000

    powerpc/eeh: Clean up EEH PEs after recovery finishes

    [ Upstream commit 799abe283e5103d48e079149579b4f167c95ea0e ]

    When the last device in an eeh_pe is removed the eeh_pe structure itself
    (and any empty parents) are freed since they are no longer needed. This
    results in a crash when a hotplug driver is involved since the following
    may occur:

    1. Device is suprise removed.
    2. Driver performs an MMIO, which fails and queues and eeh_event.
    3. Hotplug driver receives a hotplug interrupt and removes any
       pci_devs that were under the slot.
    4. pci_dev is torn down and the eeh_pe is freed.
    5. The EEH event handler thread processes the eeh_event and crashes
       since the eeh_pe pointer in the eeh_event structure is no
       longer valid.

    Crashing is generally considered poor form. Instead of doing that use
    the fact PEs are marked as EEH_PE_INVALID to keep them around until the
    end of the recovery cycle, at which point we can safely prune any empty
    PEs.

    Signed-off-by: Oliver O'Halloran <oohall@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190903101605.2890-2-oohall@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d752931fee4ff50e46ee6b5d5927c41425dc18bb
Author: Deepa Dinamani <deepa.kernel@gmail.com>
Date:   Sun Jun 23 16:00:01 2019 -0700

    pstore: fs superblock limits

    [ Upstream commit 83b8a3fbe3aa82ac3c253b698ae6a9be2dbdd5e0 ]
```

```
    Leaving granularity at 1ns because it is dependent on the specific
    attached backing pstore module. ramoops has microsecond resolution.

    Fix the readback of ramoops fractional timestamp microseconds,
    which has incorrectly been reporting the value as nanoseconds.

    Fixes: 3f8f80f0cfeb ("pstore/ram: Read and write to the 'compressed' flag of pstore").

    Signed-off-by: Deepa Dinamani <deepa.kernel@gmail.com>
    Acked-by: Kees Cook <keescook@chromium.org>
    Acked-by: Jeff Layton <jlayton@kernel.org>
    Cc: anton@enomsg.org
    Cc: ccross@android.com
    Cc: keescook@chromium.org
    Cc: tony.luck@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1154d3ce66eb01d48f9ef79ac3635ecb82999013
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Fri Aug 2 20:56:32 2019 +1000

    powerpc/64s/exception: machine check use correct cfar for late handler

    [ Upstream commit 0b66370c61fcf5fcc1d6901013e110284da6e2bb ]

    Bare metal machine checks run an "early" handler in real mode before
    running the main handler which reports the event.

    The main handler runs exactly as a normal interrupt handler, after the
    "windup" which sets registers back as they were at interrupt entry.
    CFAR does not get restored by the windup code, so that will be wrong
    when the handler is run.

    Restore the CFAR to the saved value before running the late handler.

    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190802105709.27696-8-npiggin@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84efc596dde081fa90731ae38f3b76647caed2c2
Author: Jean Delvare <jdelvare@suse.de>
Date:   Wed Aug 28 17:05:57 2019 +0200

    drm/amdgpu/si: fix ASIC tests

    [ Upstream commit 77efe48a729588527afb4d5811b9e0acb29f5e51 ]

    Comparing adev->family with CHIP constants is not correct.
    adev->family can only be compared with AMDGPU_FAMILY constants and
    adev->asic_type is the struct member to compare with CHIP constants.
    They are separate identification spaces.

    Signed-off-by: Jean Delvare <jdelvare@suse.de>
    Fixes: 62a37553414a ("drm/amdgpu: add si implementation v10")
    Cc: Ken Wang <Qingqing.Wang@amd.com>
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Cc: "Christian König" <christian.koenig@amd.com>
    Cc: "David (ChunMing) Zhou" <David1.Zhou@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 455ef08c4a87e598b8bc44890d9d720556734305
Author: Nathan Chancellor <natechancellor@gmail.com>
```

```
Date:   Mon Aug 26 17:41:55 2019 -0700

    kbuild: Do not enable -Wimplicit-fallthrough for clang for now

    [ Upstream commit e2079e93f562c7f7a030eb7642017ee5eabaaa10 ]

    This functionally reverts commit bfd77145f35c ("Makefile: Convert
    -Wimplicit-fallthrough=3 to just -Wimplicit-fallthrough for clang").

    clang enabled support for -Wimplicit-fallthrough in C in r369414 [1],
    which causes a lot of warnings when building the kernel for two reasons:

    1. Clang does not support the /* fall through */ comments. There seems
       to be a general consensus in the LLVM community that this is not
       something they want to support. Joe Perches wrote a script to convert
       all of the comments to a "fallthrough" keyword that will be added to
       compiler_attributes.h [2] [3], which catches the vast majority of the
       comments. There doesn't appear to be any consensus in the kernel
       community when to do this conversion.

    2. Clang and GCC disagree about falling through to final case statements
       with no content or cases that simply break:

       https://godbolt.org/z/c8csDu

       This difference contributes at least 50 warnings in an allyesconfig
       build for x86, not considering other architectures. This difference
       will need to be discussed to see which compiler is right [4] [5].

    [1]: https://github.com/llvm/llvm-
project/commit/1e0affb6e564b7361b0aadb38805f26deff4ecfc
    [2]:
https://lore.kernel.org/lkml/61ddbb86d5e68a15e24ccb06d9b399bbf5ce2da7.camel@perches.com/
    [3]:
https://lore.kernel.org/lkml/1d2830aadbe9d8151728a7df5b88528fc72a0095.1564549413.git.joe@per
ches.com/
    [4]: https://gcc.gnu.org/bugzilla/show_bug.cgi?id=91432
    [5]: https://github.com/ClangBuiltLinux/linux/issues/636

    Given these two problems need discussion and coordination, do not enable
    -Wimplicit-fallthrough with clang right now. Add a comment to explain
    what is going on as well. This commit should be reverted once these two
    issues are fully flushed out and resolved.

    Suggested-by: Masahiro Yamada <yamada.masahiro@socionext.com>
    Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
    Acked-by: Miguel Ojeda <miguel.ojeda.sandonis@gmail.com>
    Acked-by: Nick Desaulniers <ndesaulniers@google.com>
    Acked-by: Gustavo A. R. Silva <gustavo@embeddedor.com>
    Signed-off-by: Masahiro Yamada <yamada.masahiro@socionext.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47d112b3d9bd32f8d006f661b67bdc00db30c8ec
Author: Gustavo Romero <gromero@linux.vnet.ibm.com>
Date:   Tue Aug 20 18:54:11 2019 -0400

    selftests/powerpc: Retry on host facility unavailable

    [ Upstream commit 6652bf6408895b09d31fd4128a1589a1a0672823 ]

    TM test tm-unavailable must take into account aborts due to host aborting
    a transactin because of a facility unavailable exception, just like it
    already does for aborts on reschedules (TM_CAUSE_KVM_RESCHED).

    Reported-by: Desnes A. Nunes do Rosario <desnesn@linux.ibm.com>
```

```
    Tested-by: Desnes A. Nunes do Rosario <desnesn@linux.ibm.com>
    Signed-off-by: Gustavo Romero <gromero@linux.vnet.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/1566341651-19747-1-git-send-email-
gromero@linux.vnet.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc9c144be777530885c3b9f4ad817eecd8cd87d2
Author: Yogesh Mohan Marimuthu <yogesh.mohanmarimuthu@amd.com>
Date:   Fri Aug 2 10:52:49 2019 +0530

    drm/amd/display: fix trigger not generated for freesync

    [ Upstream commit 1e7f100ce8c0640634b794604880d9204480c9f1 ]

    [Why]
    In newer hardware MANUAL_FLOW_CONTROL is not a trigger bit. Due to this
    front porch is fixed and in these hardware freesync does not work.

    [How]
    Change the programming to generate a pulse so that the event will be
    triggered, front porch will be cut short and freesync will work.

    Signed-off-by: Yogesh Mohan Marimuthu <yogesh.mohanmarimuthu@amd.com>
    Reviewed-by: Anthony Koo <Anthony.Koo@amd.com>
    Acked-by: Bhawanpreet Lakha <Bhawanpreet.Lakha@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19a64792239cdd29078a120659aaac70c4d48e3f
Author: Zi Yu Liao <ziyu.liao@amd.com>
Date:   Tue Jul 30 15:36:53 2019 -0400

    drm/amd/display: fix MPO HUBP underflow with Scatter Gather

    [ Upstream commit 89cb5614736b9b5d3b833ca2237d10da6b4b0395 ]

    [why]
    With Scatter Gather enabled, HUBP underflows during MPO enabled video
    playback. hubp_init has a register write that fixes this problem, but
    the register is cleared when HUBP gets power gated.

    [how]
    Make a call to hubp_init during enable_plane, so that the fix can
    be applied after HUBP powers back on again.

    Signed-off-by: Zi Yu Liao <ziyu.liao@amd.com>
    Reviewed-by: Tony Cheng <Tony.Cheng@amd.com>
    Acked-by: Bhawanpreet Lakha <Bhawanpreet.Lakha@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit adaf6a48e85f91ab4b2da7848fb781fd91268f79
Author: Kevin Wang <kevin1.wang@amd.com>
Date:   Wed Aug 21 10:58:19 2019 +0800

    drm/amd/powerpaly: fix navi series custom peak level value error

    [ Upstream commit 706feb26f890e1b8297b5d14975160de361edf4f ]

    fix other navi asic set peak performance level error.
    because the navi10_ppt.c will handle navi12 14 asic,
    it will use navi10 peak value to set other asic, it is not correct.

    after patch:
```

```
    only navi10 use custom peak value, other asic will used default value.

    Signed-off-by: Kevin Wang <kevin1.wang@amd.com>
    Reviewed-by: Evan Quan <evan.quan@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d56ed01ebbd8d5b7639eeac828d47ef465ed1de9
Author: Charlene Liu <charlene.liu@amd.com>
Date:   Fri Aug 2 14:49:58 2019 -0400

    drm/amd/display: support spdif

    [ Upstream commit b5a41620bb88efb9fb31a4fa5e652e3d5bead7d4 ]

    [Description]
    port spdif fix to staging:
     spdif hardwired to afmt inst 1.
     spdif func pointer
     spdif resource allocation (reserve last audio endpoint for spdif only)

    Signed-off-by: Charlene Liu <charlene.liu@amd.com>
    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Bhawanpreet Lakha <Bhawanpreet.Lakha@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ccd23a7b78a5f9be0fadccef91d203a93ffd5a03
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Aug 9 15:44:51 2019 +0200

    clk: renesas: cpg-mssr: Set GENPD_FLAG_ALWAYS_ON for clock domain

    [ Upstream commit f787216f33ce5b5a2567766398f44ab62157114c ]

    The CPG/MSSR Clock Domain driver does not implement the
    generic_pm_domain.power_{on,off}() callbacks, as the domain itself
    cannot be powered down.  Hence the domain should be marked as always-on
    by setting the GENPD_FLAG_ALWAYS_ON flag, to prevent the core PM Domain
    code from considering it for power-off, and doing unnessary processing.

    Note that this only affects RZ/A2 SoCs.  On R-Car Gen2 and Gen3 SoCs,
    the R-Car SYSC driver handles Clock Domain creation, and offloads only
    device attachment/detachment to the CPG/MSSR driver.

    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Simon Horman <horms+renesas@verge.net.au>
    Reviewed-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48b6ddaa0ec0a3b865826bc3b4e7f22bf5eb8f75
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Aug 9 15:38:34 2019 +0200

    clk: renesas: mstp: Set GENPD_FLAG_ALWAYS_ON for clock domain

    [ Upstream commit a459a184c978ca9ad538aab93aafdde873953f30 ]

    The CPG/MSTP Clock Domain driver does not implement the
    generic_pm_domain.power_{on,off}() callbacks, as the domain itself
    cannot be powered down.  Hence the domain should be marked as always-on
    by setting the GENPD_FLAG_ALWAYS_ON flag, to prevent the core PM Domain
    code from considering it for power-off, and doing unnessary processing.

    This also gets rid of a boot warning when the Clock Domain contains an
```

```
    IRQ-safe device, e.g. on RZ/A1:

        sh_mtu2 fcff0000.timer: PM domain cpg_clocks will not be powered off

    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Simon Horman <horms+renesas@verge.net.au>
    Reviewed-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit da0a6764db3ac9eae87578cd9d8a0a007ffb85b7
Author: Daniel Drake <drake@endlessm.com>
Date:   Wed Aug 14 17:05:40 2019 +0800

    pinctrl: amd: disable spurious-firing GPIO IRQs

    [ Upstream commit d21b8adbd475dba19ac2086d3306327b4a297418 ]

    When cold-booting Asus X434DA, GPIO 7 is found to be already configured
    as an interrupt, and the GPIO level is found to be in a state that
    causes the interrupt to fire.

    As soon as pinctrl-amd probes, this interrupt fires and invokes
    amd_gpio_irq_handler(). The IRQ is acked, but no GPIO-IRQ handler was
    invoked, so the GPIO level being unchanged just causes another interrupt
    to fire again immediately after.

    This results in an interrupt storm causing this platform to hang
    during boot, right after pinctrl-amd is probed.

    Detect this situation and disable the GPIO interrupt when this happens.
    This enables the affected platform to boot as normal. GPIO 7 actually is
    the I2C touchpad interrupt line, and later on, i2c-multitouch loads and
    re-enables this interrupt when it is ready to handle it.

    Instead of this approach, I considered disabling all GPIO interrupts at
    probe time, however that seems a little risky, and I also confirmed that
    Windows does not seem to have this behaviour: the same 41 GPIO IRQs are
    enabled under both Linux and Windows, which is a far larger collection
    than the GPIOs referenced by the DSDT on this platform.

    Signed-off-by: Daniel Drake <drake@endlessm.com>
    Link: https://lore.kernel.org/r/20190814090540.7152-1-drake@endlessm.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4952cf9d0205ae148feb3caef6d3f21adf4df1c4
Author: Mark Menzynski <mmenzyns@redhat.com>
Date:   Fri Aug 2 11:21:00 2019 +0200

    drm/nouveau/volt: Fix for some cards having 0 maximum voltage

    [ Upstream commit a1af2afbd244089560794c260b2d4326a86e39b6 ]

    Some, mostly Fermi, vbioses appear to have zero max voltage. That causes Nouveau to not
parse voltage entries, thus users not being able to set higher clocks.

    When changing this value Nvidia driver still appeared to ignore it, and I wasn't able to
find out why, thus the code is ignoring the value if it is zero.

    CC: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    Signed-off-by: Mark Menzynski <mmenzyns@redhat.com>
    Reviewed-by: Karol Herbst <kherbst@redhat.com>
    Signed-off-by: Ben Skeggs <bskeggs@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 7c9b8649afd1ea17936cad4ad2441c868320ccad
Author: Ben Skeggs <bskeggs@redhat.com>
Date:   Wed May 29 15:44:57 2019 +1000

    drm/nouveau/kms/tu102-: disable input lut when input is already FP16

    [ Upstream commit 1e339ab2ac3c769c1b06b9fb7d532f8495ebc56d ]

    On Turing, an input LUT is required to transform inputs in fixed-point
    formats to FP16 for the internal display pipe.  We provide an identity
    mapping whenever a window is enabled for this reason.

    HW has error checks to ensure when the input is already FP16, that the
    input LUT is also disabled.

    Signed-off-by: Ben Skeggs <bskeggs@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f61377b7a5512d7100e96e1e77493dd97d2c3325
Author: Xiaojie Yuan <xiaojie.yuan@amd.com>
Date:   Wed Aug 21 21:00:29 2019 +0800

    drm/amdgpu/sdma5: fix number of sdma5 trap irq types for navi1x

    [ Upstream commit 9e48495017342c5d445b25eedd86d6fd884a6496 ]

    v2: set num_types based on num_instances

    navi1x has 2 sdma engines but commit
    "e7b58d03b678 drm/amdgpu: reorganize sdma v4 code to support more instances"
    changes the max number of sdma irq types (AMDGPU_SDMA_IRQ_LAST) from 2 to 8
    which causes amdgpu_irq_gpu_reset_resume_helper() to recover irq of sdma
    engines with following logic:

    (enable irq for sdma0) * 1 time
    (enable irq for sdma1) * 1 time
    (disable irq for sdma1) * 6 times

    as a result, after gpu reset, interrupt for sdma1 is lost.

    Signed-off-by: Xiaojie Yuan <xiaojie.yuan@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 7a310d3f722caf2d0ecf621b6b2a80fc25417ad1
Author: hexin <hexin.op@gmail.com>
Date:   Thu Aug 22 11:35:19 2019 +0800

    vfio_pci: Restore original state on release

    [ Upstream commit 92c8026854c25093946e0d7fe536fd9eac440f06 ]

    vfio_pci_enable() saves the device's initial configuration information
    with the intent that it is restored in vfio_pci_disable().  However,
    the commit referenced in Fixes: below replaced the call to
    __pci_reset_function_locked(), which is not wrapped in a state save
    and restore, with pci_try_reset_function(), which overwrites the
    restored device state with the current state before applying it to the
    device.  Reinstate use of __pci_reset_function_locked() to return to
    the desired behavior.

    Fixes: 890ed578df82 ("vfio-pci: Use pci "try" reset interface")
    Signed-off-by: hexin <hexin15@baidu.com>
    Signed-off-by: Liu Qi <liuqi16@baidu.com>
```

```
    Signed-off-by: Zhang Yu <zhangyu31@baidu.com>
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 55cc62e9718f256a783bb6479d5ea346444cc22f
Author: Sam Bobroff <sbobroff@linux.ibm.com>
Date:   Fri Aug 16 14:48:06 2019 +1000

    powerpc/eeh: Clear stale EEH_DEV_NO_HANDLER flag

    [ Upstream commit aa06e3d60e245284d1e55497eb3108828092818d ]

    The EEH_DEV_NO_HANDLER flag is used by the EEH system to prevent the
    use of driver callbacks in drivers that have been bound part way
    through the recovery process. This is necessary to prevent later stage
    handlers from being called when the earlier stage handlers haven't,
    which can be confusing for drivers.

    However, the flag is set for all devices that are added after boot
    time and only cleared at the end of the EEH recovery process. This
    results in hot plugged devices erroneously having the flag set during
    the first recovery after they are added (causing their driver's
    handlers to be incorrectly ignored).

    To remedy this, clear the flag at the beginning of recovery
    processing. The flag is still cleared at the end of recovery
    processing, although it is no longer really necessary.

    Also clear the flag during eeh_handle_special_event(), for the same
    reasons.

    Signed-off-by: Sam Bobroff <sbobroff@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/b8ca5629d27de74c957d4f4b250177d1b6fc4bbd.1565930772.git.sbobroff@l
inux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 942e07ec6eb093911e5942171ac14c3afbac9768
Author: Sowjanya Komatineni <skomatineni@nvidia.com>
Date:   Fri Aug 16 12:41:46 2019 -0700

    pinctrl: tegra: Fix write barrier placement in pmx_writel

    [ Upstream commit c2cf351eba2ff6002ce8eb178452219d2521e38e ]

    pmx_writel uses writel which inserts write barrier before the
    register write.

    This patch has fix to replace writel with writel_relaxed followed
    by a readback and memory barrier to ensure write operation is
    completed for successful pinctrl change.

    Acked-by: Thierry Reding <treding@nvidia.com>
    Reviewed-by: Dmitry Osipenko <digetx@gmail.com>
    Signed-off-by: Sowjanya Komatineni <skomatineni@nvidia.com>
    Link: https://lore.kernel.org/r/1565984527-5272-2-git-send-email-skomatineni@nvidia.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 74849ab2fa0c256c576c68507b61f5c4b1ff8340
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Wed Jul 24 18:46:36 2019 +1000

    powerpc/perf: fix imc allocation failure handling
```

```
    [ Upstream commit 10c4bd7cd28e77aeb8cfa65b23cb3c632ede2a49 ]

    The alloc_pages_node return value should be tested for failure
    before being passed to page_address.

    Tested-by: Anju T Sudhakar <anju@linux.vnet.ibm.com>
    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Reviewed-by: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190724084638.24982-3-npiggin@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d352ce9a1af4c2d756b48f47e1a2a5a3f43564f3
Author: Nathan Lynch <nathanl@linux.ibm.com>
Date:   Fri Aug 2 14:29:26 2019 -0500

    powerpc/pseries/mobility: use cond_resched when updating device tree

    [ Upstream commit ccfb5bd71d3d1228090a8633800ae7cdf42a94ac ]

    After a partition migration, pseries_devicetree_update() processes
    changes to the device tree communicated from the platform to
    Linux. This is a relatively heavyweight operation, with multiple
    device tree searches, memory allocations, and conversations with
    partition firmware.

    There's a few levels of nested loops which are bounded only by
    decisions made by the platform, outside of Linux's control, and indeed
    we have seen RCU stalls on large systems while executing this call
    graph. Use cond_resched() in these loops so that the cpu is yielded
    when needed.

    Signed-off-by: Nathan Lynch <nathanl@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190802192926.19277-4-nathanl@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad8c26bf3020e8e979bac4271059c5e1ec2ea645
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Wed Jul 24 18:46:34 2019 +1000

    powerpc/64s/radix: Fix memory hotplug section page table creation

    [ Upstream commit 8f51e3929470942e6a8744061254fdeef646cd36 ]

    create_physical_mapping expects physical addresses, but creating and
    splitting these mappings after boot is supplying virtual (effective)
    addresses. This can be irritated by booting with mem= to limit memory
    then probing an unused physical memory range:

      echo <addr> > /sys/devices/system/memory/probe

    This mostly works by accident, firstly because __va(__va(x)) == __va(x)
    so the virtual address does not get corrupted. Secondly because pfn_pte
    masks out the upper bits of the pfn beyond the physical address limit,
    so a pfn constructed with a 0xc000000000000000 virtual linear address
    will be masked back to the correct physical address in the pte.

    Fixes: 6cc27341b21a8 ("powerpc/mm: add radix__create_section_mapping()")
    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Reviewed-by: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190724084638.24982-1-npiggin@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 57d9067c0fab21370238e892213d044e718daa85
Author: Christophe Leroy <christophe.leroy@c-s.fr>
Date:   Wed Aug 14 09:25:52 2019 +0000

    powerpc/futex: Fix warning: 'oldval' may be used uninitialized in this function

    [ Upstream commit 38a0d0cdb46d3f91534e5b9839ec2d67be14c59d ]

    We see warnings such as:
      kernel/futex.c: In function 'do_futex':
      kernel/futex.c:1676:17: warning: 'oldval' may be used uninitialized in this function
[-Wmaybe-uninitialized]
         return oldval == cmparg;
                ^
      kernel/futex.c:1651:6: note: 'oldval' was declared here
        int oldval, ret;
            ^

    This is because arch_futex_atomic_op_inuser() only sets *oval if ret
    is 0 and GCC doesn't see that it will only use it when ret is 0.

    Anyway, the non-zero ret path is an error path that won't suffer from
    setting *oval, and as *oval is a local var in futex_atomic_op_inuser()
    it will have no impact.

    Signed-off-by: Christophe Leroy <christophe.leroy@c-s.fr>
    [mpe: reword change log slightly]
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/86b72f0c134367b214910b27b9a6dd3321af93bb.1565774657.git.christophe
.leroy@c-s.fr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5a82cb13541bf602f9ddf6963b4ceb55ccbc65b
Author: Christophe Leroy <christophe.leroy@c-s.fr>
Date:   Wed Aug 14 12:36:10 2019 +0000

    powerpc/ptdump: fix walk_pagetables() address mismatch

    [ Upstream commit e033829d2aaad85cb7cf46986c3be0bcc72f791e ]

    walk_pagetables() always walk the entire pgdir from address 0
    but considers PAGE_OFFSET or KERN_VIRT_START as the starting
    address of the walk, resulting in a possible mismatch in the
    displayed addresses.

    Ex: on PPC32, when KERN_VIRT_START was locally defined as
    PAGE_OFFSET, ptdump displayed 0x80000000
    instead of 0xc0000000 for the first kernel page,
    because 0xc0000000 + 0xc0000000 = 0x80000000

    Start the walk at st->start_address instead of starting at 0.

    Signed-off-by: Christophe Leroy <christophe.leroy@c-s.fr>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/5aa2ac513295f594cce8ddb1c649f61947bd063d.1565786091.git.christophe
.leroy@c-s.fr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 288c27328caf5496d6d66d7c73f7afa488a37048
Author: Nathan Lynch <nathanl@linux.ibm.com>
Date:   Fri Aug 2 14:29:24 2019 -0500
```

```
    powerpc/rtas: use device model APIs and serialization during LPM

    [ Upstream commit a6717c01ddc259f6f73364779df058e2c67309f8 ]

    The LPAR migration implementation and userspace-initiated cpu hotplug
    can interleave their executions like so:

    1. Set cpu 7 offline via sysfs.

    2. Begin a partition migration, whose implementation requires the OS
       to ensure all present cpus are online; cpu 7 is onlined:

         rtas_ibm_suspend_me -> rtas_online_cpus_mask -> cpu_up

       This sets cpu 7 online in all respects except for the cpu's
       corresponding struct device; dev->offline remains true.

    3. Set cpu 7 online via sysfs. _cpu_up() determines that cpu 7 is
       already online and returns success. The driver core (device_online)
       sets dev->offline = false.

    4. The migration completes and restores cpu 7 to offline state:

         rtas_ibm_suspend_me -> rtas_offline_cpus_mask -> cpu_down

    This leaves cpu7 in a state where the driver core considers the cpu
    device online, but in all other respects it is offline and
    unused. Attempts to online the cpu via sysfs appear to succeed but the
    driver core actually does not pass the request to the lower-level
    cpuhp support code. This makes the cpu unusable until the cpu device
    is manually set offline and then online again via sysfs.

    Instead of directly calling cpu_up/cpu_down, the migration code should
    use the higher-level device core APIs to maintain consistent state and
    serialize operations.

    Fixes: 120496ac2d2d ("powerpc: Bring all threads online prior to migration/hibernation")
    Signed-off-by: Nathan Lynch <nathanl@linux.ibm.com>
    Reviewed-by: Gautham R. Shenoy <ego@linux.vnet.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190802192926.19277-2-nathanl@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbcd9badd30de7473e8af71150419b21c0d5be1e
Author: Cédric Le Goater <clg@kaod.org>
Date:   Wed Aug 14 17:47:52 2019 +0200

    powerpc/xmon: Check for HV mode when dumping XIVE info from OPAL

    [ Upstream commit c3e0dbd7f780a58c4695f1cd8fc8afde80376737 ]

    Currently, the xmon 'dx' command calls OPAL to dump the XIVE state in
    the OPAL logs and also outputs some of the fields of the internal XIVE
    structures in Linux. The OPAL calls can only be done on baremetal
    (PowerNV) and they crash a pseries machine. Fix by checking the
    hypervisor feature of the CPU.

    Signed-off-by: Cédric Le Goater <clg@kaod.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190814154754.23682-2-clg@kaod.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a8a92561e41ac7b0caf9fb2e71779784811b719e
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Wed Aug 14 21:10:37 2019 -0700
```

```
    clk: sunxi: Don't call clk_hw_get_name() on a hw that isn't registered

    [ Upstream commit a7b85ad25a97cf897b4819a121655c483d86156f ]

    The implementation of clk_hw_get_name() relies on the clk_core
    associated with the clk_hw pointer existing. If of_clk_hw_register()
    fails, there isn't a clk_core created yet, so calling clk_hw_get_name()
    here fails. Extract the name first so we can print it later.

    Fixes: 1d80c14248d6 ("clk: sunxi-ng: Add common infrastructure")
    Cc: Maxime Ripard <maxime.ripard@bootlin.com>
    Cc: Chen-Yu Tsai <wens@csie.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b63620cdcb408ac91aa95574e9de937437bbc00
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Thu Aug 15 09:00:18 2019 -0700

    clk: zx296718: Don't reference clk_init_data after registration

    [ Upstream commit 1a4549c150e27dbc3aea762e879a88209df6d1a5 ]

    A future patch is going to change semantics of clk_register() so that
    clk_hw::init is guaranteed to be NULL after a clk is registered. Avoid
    referencing this member here so that we don't run into NULL pointer
    exceptions.

    Cc: Jun Nie <jun.nie@linaro.org>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lkml.kernel.org/r/20190815160020.183334-3-sboyd@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e873712459e60a679cc111f7b153f2bd619c503a
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Wed Jul 31 12:35:15 2019 -0700

    clk: sprd: Don't reference clk_init_data after registration

    [ Upstream commit f6c90df8e7e33c3dc33d4d7471bc42c232b0510e ]

    A future patch is going to change semantics of clk_register() so that
    clk_hw::init is guaranteed to be NULL after a clk is registered. Avoid
    referencing this member here so that we don't run into NULL pointer
    exceptions.

    Cc: Chunyan Zhang <zhang.chunyan@linaro.org>
    Cc: Baolin Wang <baolin.wang@linaro.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lkml.kernel.org/r/20190731193517.237136-8-sboyd@kernel.org
    Acked-by: Baolin Wang <baolin.wang@linaro.org>
    Acked-by: Chunyan Zhang <zhang.chunyan@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba44fed3404448fd9b8baea72c66d72db9ba5d3a
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Wed Jul 31 12:35:11 2019 -0700

    clk: meson: axg-audio: Don't reference clk_init_data after registration

    [ Upstream commit 1610dd79d0f6202c5c1a91122255fa598679c13a ]

    A future patch is going to change semantics of clk_register() so that
```

```
    clk_hw::init is guaranteed to be NULL after a clk is registered. Avoid
    referencing this member here so that we don't run into NULL pointer
    exceptions.

    Cc: Neil Armstrong <narmstrong@baylibre.com>
    Cc: Jerome Brunet <jbrunet@baylibre.com>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lkml.kernel.org/r/20190731193517.237136-4-sboyd@kernel.org
    Acked-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 532a30d3aaeff4c9f69f27f016edc3f47639dba4
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Wed Jul 31 12:35:13 2019 -0700

    clk: sirf: Don't reference clk_init_data after registration

    [ Upstream commit af55dadfbce35b4f4c6247244ce3e44b2e242b84 ]

    A future patch is going to change semantics of clk_register() so that
    clk_hw::init is guaranteed to be NULL after a clk is registered. Avoid
    referencing this member here so that we don't run into NULL pointer
    exceptions.

    Cc: Guo Zeng <Guo.Zeng@csr.com>
    Cc: Barry Song <Baohua.Song@csr.com>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lkml.kernel.org/r/20190731193517.237136-6-sboyd@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f72e322328e2549cd19c66b26a5d594844c70741
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Wed Jul 31 12:35:09 2019 -0700

    clk: actions: Don't reference clk_init_data after registration

    [ Upstream commit cf9ec1fc6d7cceb73e7f1efd079d2eae173fdf57 ]

    A future patch is going to change semantics of clk_register() so that
    clk_hw::init is guaranteed to be NULL after a clk is registered. Avoid
    referencing this member here so that we don't run into NULL pointer
    exceptions.

    Cc: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lkml.kernel.org/r/20190731193517.237136-2-sboyd@kernel.org
    [sboyd@kernel.org: Move name to after checking for error or NULL hw]
    Acked-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e14d52b0234a1d5240db49aa2d3b9150f161e252
Author: Alexey Kardashevskiy <aik@ozlabs.ru>
Date:   Thu Jul 18 15:11:38 2019 +1000

    powerpc/powernv/ioda2: Allocate TCE table levels on demand for default DMA window

    [ Upstream commit c37c792dec0929dbb6360a609fb00fa20bb16fc2 ]

    We allocate only the first level of multilevel TCE tables for KVM
    already (alloc_userspace_copy==true), and the rest is allocated on demand.
    This is not enabled though for bare metal.

    This removes the KVM limitation (implicit, via the alloc_userspace_copy
    parameter) and always allocates just the first level. The on-demand
    allocation of missing levels is already implemented.
```

```
    As from now on DMA map might happen with disabled interrupts, this
    allocates TCEs with GFP_ATOMIC; otherwise lockdep reports errors 1].
    In practice just a single page is allocated there so chances for failure
    are quite low.

    To save time when creating a new clean table, this skips non-allocated
    indirect TCE entries in pnv_tce_free just like we already do in
    the VFIO IOMMU TCE driver.

    This changes the default level number from 1 to 2 to reduce the amount
    of memory required for the default 32bit DMA window at the boot time.
    The default window size is up to 2GB which requires 4MB of TCEs which is
    unlikely to be used entirely or at all as most devices these days are
    64bit capable so by switching to 2 levels by default we save 4032KB of
    RAM per a device.

    While at this, add __GFP_NOWARN to alloc_pages_node() as the userspace
    can trigger this path via VFIO, see the failure and try creating a table
    again with different parameters which might succeed.

    [1]:
    ===
    BUG: sleeping function called from invalid context at mm/page_alloc.c:4596
    in_atomic(): 1, irqs_disabled(): 1, pid: 1038, name: scsi_eh_1
    2 locks held by scsi_eh_1/1038:
     #0: 000000005efd659a (&host->eh_mutex){+.+.}, at: ata_eh_acquire+0x34/0x80
     #1: 0000000006cf56a6 (&(&host->lock)->rlock){....}, at: ata_exec_internal_sg+0xb0/0x5c0
    irq event stamp: 500
    hardirqs last  enabled at (499): [<c000000000cb8a74>]
_raw_spin_unlock_irqrestore+0x94/0xd0
    hardirqs last disabled at (500): [<c000000000cb85c4>] _raw_spin_lock_irqsave+0x44/0x120
    softirqs last  enabled at (0): [<c000000000101120>]
copy_process.isra.4.part.5+0x640/0x1a80
    softirqs last disabled at (0): [<0000000000000000>] 0x0
    CPU: 73 PID: 1038 Comm: scsi_eh_1 Not tainted 5.2.0-rc6-le_nv2_aikATfstn1-p1 #634
    Call Trace:
    [c000003d064cef50] [c000000000c8e6c4] dump_stack+0xe8/0x164 (unreliable)
    [c000003d064cefa0] [c00000000014ed78] ___might_sleep+0x2f8/0x310
    [c000003d064cf020] [c0000000003ca084] __alloc_pages_nodemask+0x2a4/0x1560
    [c000003d064cf220] [c0000000000c2530] pnv_alloc_tce_level.isra.0+0x90/0x130
    [c000003d064cf290] [c0000000000c2888] pnv_tce+0x128/0x3b0
    [c000003d064cf360] [c0000000000c2c00] pnv_tce_build+0xb0/0xf0
    [c000003d064cf3c0] [c0000000000bbd9c] pnv_ioda2_tce_build+0x3c/0xb0
    [c000003d064cf400] [c00000000004cfe0] ppc_iommu_map_sg+0x210/0x550
    [c000003d064cf510] [c00000000004b7a4] dma_iommu_map_sg+0x74/0xb0
    [c000003d064cf530] [c000000000863944] ata_qc_issue+0x134/0x470
    [c000003d064cf5b0] [c000000000863ec4] ata_exec_internal_sg+0x244/0x5c0
    [c000003d064cf700] [c0000000008642d0] ata_exec_internal+0x90/0xe0
    [c000003d064cf780] [c0000000008650ac] ata_dev_read_id+0x2ec/0x640
    [c000003d064cf8d0] [c000000000878e28] ata_eh_recover+0x948/0x16d0
    [c000003d064cfa10] [c00000000087d760] sata_pmp_error_handler+0x480/0xbf0
    [c000003d064cfbc0] [c000000000884624] ahci_error_handler+0x74/0xe0
    [c000003d064cfbf0] [c000000000879fa8] ata_scsi_port_error_handler+0x2d8/0x7c0
    [c000003d064cfca0] [c00000000087a544] ata_scsi_error+0xb4/0x100
    [c000003d064cfd00] [c000000000802450] scsi_error_handler+0x120/0x510
    [c000003d064cfdb0] [c000000000140c48] kthread+0x1b8/0x1c0
    [c000003d064cfe20] [c00000000000bd8c] ret_from_kernel_thread+0x5c/0x70
    ata1: SATA link up 6.0 Gbps (SStatus 133 SControl 300)
    irq event stamp: 2305

    ======================================================
    hardirqs last  enabled at (2305): [<c00000000000e4c8>] fast_exc_return_irq+0x28/0x34
    hardirqs last disabled at (2303): [<c000000000cb9fd0>] __do_softirq+0x4a0/0x654
    WARNING: possible irq lock inversion dependency detected
```

```
     5.2.0-rc6-le_nv2_aikATfstn1-p1 #634 Tainted: G        W
     softirqs last  enabled at (2304): [<c000000000cba054>] __do_softirq+0x524/0x654
     softirqs last disabled at (2297): [<c00000000010f278>] irq_exit+0x128/0x180
     -----------------------------------------------------
     swapper/0/0 just changed the state of lock:
     0000000006cf56a6 (&(&host->lock)->rlock){-...}, at:
  ahci_single_level_irq_intr+0xac/0x120
     but this lock took another, HARDIRQ-unsafe lock in the past:
      (fs_reclaim){+.+.}

     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(fs_reclaim);
                                    local_irq_disable();
                                    lock(&(&host->lock)->rlock);
                                    lock(fs_reclaim);
       <Interrupt>
         lock(&(&host->lock)->rlock);

      *** DEADLOCK ***

     no locks held by swapper/0/0.

     the shortest dependencies between 2nd lock and 1st lock:
      -> (fs_reclaim){+.+.} ops: 167579 {
         HARDIRQ-ON-W at:
                             lock_acquire+0xf8/0x2a0
                             fs_reclaim_acquire.part.23+0x44/0x60
                             kmem_cache_alloc_node_trace+0x80/0x590
                             alloc_desc+0x64/0x270
                             __irq_alloc_descs+0x2e4/0x3a0
                             irq_domain_alloc_descs+0xb0/0x150
                             irq_create_mapping+0x168/0x2c0
                             xics_smp_probe+0x2c/0x98
                             pnv_smp_probe+0x40/0x9c
                             smp_prepare_cpus+0x524/0x6c4
                             kernel_init_freeable+0x1b4/0x650
                             kernel_init+0x2c/0x148
                             ret_from_kernel_thread+0x5c/0x70
         SOFTIRQ-ON-W at:
                             lock_acquire+0xf8/0x2a0
                             fs_reclaim_acquire.part.23+0x44/0x60
                             kmem_cache_alloc_node_trace+0x80/0x590
                             alloc_desc+0x64/0x270
                             __irq_alloc_descs+0x2e4/0x3a0
                             irq_domain_alloc_descs+0xb0/0x150
                             irq_create_mapping+0x168/0x2c0
                             xics_smp_probe+0x2c/0x98
                             pnv_smp_probe+0x40/0x9c
                             smp_prepare_cpus+0x524/0x6c4
                             kernel_init_freeable+0x1b4/0x650
                             kernel_init+0x2c/0x148
                             ret_from_kernel_thread+0x5c/0x70
         INITIAL USE at:
                            lock_acquire+0xf8/0x2a0
                            fs_reclaim_acquire.part.23+0x44/0x60
                            kmem_cache_alloc_node_trace+0x80/0x590
                            alloc_desc+0x64/0x270
                            __irq_alloc_descs+0x2e4/0x3a0
                            irq_domain_alloc_descs+0xb0/0x150
```

```
                        irq_create_mapping+0x168/0x2c0
                        xics_smp_probe+0x2c/0x98
                        pnv_smp_probe+0x40/0x9c
                        smp_prepare_cpus+0x524/0x6c4
                        kernel_init_freeable+0x1b4/0x650
                        kernel_init+0x2c/0x148
                        ret_from_kernel_thread+0x5c/0x70
      }
    ===

    Signed-off-by: Alexey Kardashevskiy <aik@ozlabs.ru>
    Reviewed-by: Alistair Popple <alistair@popple.id.au>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20190718051139.74787-4-aik@ozlabs.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff9db28c593711a16496850603e98c4d0f8f0239
Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
Date:   Thu Jul 25 11:53:16 2019 -0400

    drm/amd/display: Register VUPDATE_NO_LOCK interrupts for DCN2

    [ Upstream commit e40837afb9b011757e17e9f71d97853ca574bcff ]

    [Why]
    These are needed to send back DRM vblank events in the case where VRR
    is on. Without the interrupt enabled we're deferring the events into the
    vblank queue and userspace is left waiting forever to get back the
    events they need.

    Found using igt@kms_vrr - the test fails immediately due to vblank
    timeout.

    [How]
    Register them the same way we're handling it for DCN1.

    This fixes igt@kms_vrr for DCN2.

    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Reviewed-by: David Francis <David.Francis@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4160705dcfcf2cbd5f17caee872d1b2bbee6ac6b
Author: Lewis Huang <Lewis.Huang@amd.com>
Date:   Mon Jul 22 15:23:32 2019 -0400

    drm/amd/display: reprogram VM config when system resume

    [ Upstream commit e5382701c3520b3ed66169a6e4aa6ce5df8c56e0 ]

    [Why]
    The vm config will be clear to 0 when system enter S4. It will
    cause hubbub didn't know how to fetch data when system resume.
    The flip always pending because earliest_inuse_address and
    request_address are different.

    [How]
    Reprogram VM config when system resume

    Signed-off-by: Lewis Huang <Lewis.Huang@amd.com>
    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Eric Yang <eric.yang2@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
```

```
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eca427e6494005cb10d7fc24e57d870033a0ce4b
Author: Bayan Zabihiyan <bayan.zabihiyan@amd.com>
Date:   Wed Jul 10 16:00:53 2019 -0400

    drm/amd/display: Fix frames_to_insert math

    [ Upstream commit a463b263032f7c98c5912207db43be1aa34a6438 ]

    [Why]
    The math on deciding on how many
    "frames to insert" sometimes sent us over the max refresh rate.
    Also integer overflow can occur if we have high refresh rates.

    [How]
    Instead of clipping the  frame duration such that it doesn’t go below the min,
    just remove a frame from the number of frames to insert. +
    Use unsigned long long for intermediate calculations to prevent
    integer overflow.

    Signed-off-by: Bayan Zabihiyan <bayan.zabihiyan@amd.com>
    Reviewed-by: Aric Cyr <Aric.Cyr@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e7a1d09a5c93d086c3b2687e43a80285cd198d5
Author: Anthony Koo <Anthony.Koo@amd.com>
Date:   Fri Jul 12 10:52:54 2019 -0400

    drm/amd/display: fix issue where 252-255 values are clipped

    [ Upstream commit 1cbcfc975164f397b449efb17f59d81a703090db ]

    [Why]
    When endpoint is at the boundary of a region, such as at 2^0=1
    we find that the last segment has a sharp slope and some points
    are clipped at the top.

    [How]
    If end point is 1, which is exactly at the 2^0 region boundary, we
    need to program an additional region beyond this point.

    Signed-off-by: Anthony Koo <Anthony.Koo@amd.com>
    Reviewed-by: Aric Cyr <Aric.Cyr@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 53bb952338564a4b06f0234bc7ca05a08a5be101
Author: Icenowy Zheng <icenowy@aosc.io>
Date:   Sun Jul 28 11:12:23 2019 +0800

    clk: sunxi-ng: v3s: add missing clock slices for MMC2 module clocks

    [ Upstream commit 720099603d1f62e37b789366d7e89824b009ca28 ]

    The MMC2 clock slices are currently not defined in V3s CCU driver, which
    makes MMC2 not working.

    Fix this issue.

    Fixes: d0f11d14b0bc ("clk: sunxi-ng: add support for V3s CCU")
```

```
    Signed-off-by: Icenowy Zheng <icenowy@aosc.io>
    Signed-off-by: Maxime Ripard <maxime.ripard@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c041d1dc67e5499cc734d06b001d9ea5215ccad
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Mon Jul 1 13:36:06 2019 +0200

    clk: ingenic/jz4740: Fix "pll half" divider not read/written properly

    [ Upstream commit 568b9de48d80bcf1a92e2c4fa67651abbb8ebfe2 ]

    The code was setting the bit 21 of the CPCCR register to use a divider
    of 2 for the "pll half" clock, and clearing the bit to use a divider
    of 1.

    This is the opposite of how this register field works: a cleared bit
    means that the /2 divider is used, and a set bit means that the divider
    is 1.

    Restore the correct behaviour using the newly introduced .div_table
    field.

    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Link: https://lkml.kernel.org/r/20190701113606.4130-1-paul@crapouillou.net
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50647140391d987658d7e254405fa64f2d40eee2
Author: Nathan Huckleberry <nhuck@google.com>
Date:   Thu Jun 27 15:06:42 2019 -0700

    clk: qoriq: Fix -Wunused-const-variable

    [ Upstream commit a95fb581b144b5e73da382eaedb2e32027610597 ]

    drivers/clk/clk-qoriq.c:138:38: warning: unused variable
    'p5020_cmux_grp1' [-Wunused-const-variable] static const struct
    clockgen_muxinfo p5020_cmux_grp1

    drivers/clk/clk-qoriq.c:146:38: warning: unused variable
    'p5020_cmux_grp2' [-Wunused-const-variable] static const struct
    clockgen_muxinfo p5020_cmux_grp2

    In the definition of the p5020 chip, the p2041 chip's info was used
    instead.  The p5020 and p2041 chips have different info. This is most
    likely a typo.

    Link: https://github.com/ClangBuiltLinux/linux/issues/525
    Cc: clang-built-linux@googlegroups.com
    Signed-off-by: Nathan Huckleberry <nhuck@google.com>
    Link: https://lkml.kernel.org/r/20190627220642.78575-1-nhuck@google.com
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Acked-by: Scott Wood <oss@buserror.net>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f75ba2f4522d4b5a626a2f645053e935b6b253e3
Author: Corey Minyard <cminyard@mvista.com>
Date:   Fri Aug 2 07:31:36 2019 -0500

    ipmi_si: Only schedule continuously in the thread in maintenance mode

    [ Upstream commit 340ff31ab00bca5c15915e70ad9ada3030c98cf8 ]
```

```
    ipmi_thread() uses back-to-back schedule() to poll for command
    completion which, on some machines, can push up CPU consumption and
    heavily tax the scheduler locks leading to noticeable overall
    performance degradation.

    This was originally added so firmware updates through IPMI would
    complete in a timely manner.  But we can't kill the scheduler
    locks for that one use case.

    Instead, only run schedule() continuously in maintenance mode,
    where firmware updates should run.

    Signed-off-by: Corey Minyard <cminyard@mvista.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1c41ac2054bd357f8213dbf0b0f25456d928c21e
Author: Alexandre Torgue <alexandre.torgue@st.com>
Date:   Thu Jul 25 13:16:56 2019 +0200

    pinctrl: stmfx: update pinconf settings

    [ Upstream commit a502b343ebd0eab38f3cb33fbb84011847cf5aac ]

    According to the following tab (coming from STMFX datasheet), updates
    have to done in stmfx_pinconf_set function:

    -"type" has to be set when "bias" is configured as "pull-up or pull-down"
    -PIN_CONFIG_DRIVE_PUSH_PULL should only be used when gpio is configured as
     output. There is so no need to check direction.

    DIR | TYPE | PUPD | MFX GPIO configuration
    ----|------|------|-----------------------------------------------------
    1   | 1    | 1    | OUTPUT open drain with internal pull-up resistor
    ----|------|------|-----------------------------------------------------
    1   | 1    | 0    | OUTPUT open drain with internal pull-down resistor
    ----|------|------|-----------------------------------------------------
    1   | 0    | 0/1  | OUTPUT push pull no pull
    ----|------|------|-----------------------------------------------------
    0   | 1    | 1    | INPUT with internal pull-up resistor
    ----|------|------|-----------------------------------------------------
    0   | 1    | 0    | INPUT with internal pull-down resistor
    ----|------|------|-----------------------------------------------------
    0   | 0    | 1    | INPUT floating
    ----|------|------|-----------------------------------------------------
    0   | 0    | 0    | analog (GPIO not used, default setting)

    Signed-off-by: Alexandre Torgue <alexandre.torgue@st.com>
    Signed-off-by: Amelie Delaunay <amelie.delaunay@st.com>
    Link: https://lore.kernel.org/r/1564053416-32192-1-git-send-email-amelie.delaunay@st.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df05cb10f25e63b5c4b99046bda75ce5a3d21656
Author: Nathan Chancellor <natechancellor@gmail.com>
Date:   Mon Jun 3 15:11:58 2019 -0700

    PCI: rpaphp: Avoid a sometimes-uninitialized warning

    [ Upstream commit 0df3e42167caaf9f8c7b64de3da40a459979afe8 ]

    When building with -Wsometimes-uninitialized, clang warns:

    drivers/pci/hotplug/rpaphp_core.c:243:14: warning: variable 'fndit' is
    used uninitialized whenever 'for' loop exits because its condition is
    false [-Wsometimes-uninitialized]
```

```
            for (j = 0; j < entries; j++) {
                        ^~~~~~~~~~
    drivers/pci/hotplug/rpaphp_core.c:256:6: note: uninitialized use occurs
    here
            if (fndit)
                ^~~~
    drivers/pci/hotplug/rpaphp_core.c:243:14: note: remove the condition if
    it is always true
            for (j = 0; j < entries; j++) {
                        ^~~~~~~~~~
    drivers/pci/hotplug/rpaphp_core.c:233:14: note: initialize the variable
    'fndit' to silence this warning
            int j, fndit;
                        ^
                         = 0
```

fndit is only used to gate a sprintf call, which can be moved into the loop to simplify the code and eliminate the local variable, which will fix this warning.

Fixes: 2fcf3ae508c2 ("hotplug/drc-info: Add code to search ibm,drc-info property")
Suggested-by: Nick Desaulniers <ndesaulniers@google.com>
Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
Acked-by: Tyrel Datwyler <tyreld@linux.ibm.com>
Acked-by: Joel Savitz <jsavitz@redhat.com>
Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
Link: https://github.com/ClangBuiltLinux/linux/issues/504
Link: https://lore.kernel.org/r/20190603221157.58502-1-natechancellor@gmail.com
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e62799d46ac210f34d8adbe5a16b855be4ac9a0
Author: Abel Vesa <abel.vesa@nxp.com>
Date: Tue Jul 30 10:22:55 2019 +0300

clk: imx8mq: Mark AHB clock as critical

[Upstream commit 9b9c60bed562c3718ae324a86f3f30a4ff983cf8]

Initially, the TMU_ROOT clock was marked as critical, which automatically made the AHB clock to stay always on. Since the TMU_ROOT clock is not marked as critical anymore, following commit:

"clk: imx8mq: Remove CLK_IS_CRITICAL flag for IMX8MQ_CLK_TMU_ROOT"

all the clocks that derive from ipg_root clock (and implicitly ahb clock) would also have to enable, along with their own gate, the AHB clock.

But considering that AHB is actually a bus that has to be always on, we mark it as critical in the clock provider driver and then all the clocks that derive from it can be controlled through the dedicated per IP gate which follows after the ipg_root clock.

Signed-off-by: Abel Vesa <abel.vesa@nxp.com>
Tested-by: Daniel Baluta <daniel.baluta@nxp.com>
Signed-off-by: Shawn Guo <shawnguo@kernel.org>
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d36ff44ec01923c4f9f18e3f34efeddc433ea08a
Author: Jia-Ju Bai <baijiaju1990@gmail.com>
Date: Mon Jul 29 16:36:44 2019 +0800

gpu: drm: radeon: Fix a possible null-pointer dereference in radeon_connector_set_property()

[Upstream commit f3eb9b8f67bc28783eddc142ad805ebdc53d6339]

```
    In radeon_connector_set_property(), there is an if statement on line 743
    to check whether connector->encoder is NULL:
        if (connector->encoder)

    When connector->encoder is NULL, it is used on line 755:
        if (connector->encoder->crtc)

    Thus, a possible null-pointer dereference may occur.

    To fix this bug, connector->encoder is checked before being used.

    This bug is found by a static analysis tool STCheck written by us.

    Signed-off-by: Jia-Ju Bai <baijiaju1990@gmail.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 894c414129a8d9ef1b2de443015e4dde6085f64f
Author: KyleMahlkuch <kmahlkuc@linux.vnet.ibm.com>
Date:   Wed Jul 31 17:10:14 2019 -0500

    drm/radeon: Fix EEH during kexec

    [ Upstream commit 6f7fe9a93e6c09bf988c5059403f5f88e17e21e6 ]

    During kexec some adapters hit an EEH since they are not properly
    shut down in the radeon_pci_shutdown() function. Adding
    radeon_suspend_kms() fixes this issue.

    Signed-off-by: KyleMahlkuch <kmahlkuc@linux.vnet.ibm.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58ed657ef16315c4801a6a8ca4880e8905c2b909
Author: Nathan Chancellor <natechancellor@gmail.com>
Date:   Wed Jul 3 22:52:16 2019 -0700

    drm/amd/display: Use proper enum conversion functions

    [ Upstream commit d196bbbc28fab82624f7686f8b0da8e8644b6e6a ]

    clang warns:

    drivers/gpu/drm/amd/amdgpu/../display/amdgpu_dm/amdgpu_dm_pp_smu.c:336:8:
    warning: implicit conversion from enumeration type 'enum smu_clk_type'
    to different enumeration type 'enum amd_pp_clock_type'
    [-Wenum-conversion]
                                            dc_to_smu_clock_type(clk_type),
                                            ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    drivers/gpu/drm/amd/amdgpu/../display/amdgpu_dm/amdgpu_dm_pp_smu.c:421:14:
    warning: implicit conversion from enumeration type 'enum
    amd_pp_clock_type' to different enumeration type 'enum smu_clk_type'
    [-Wenum-conversion]
                                            dc_to_pp_clock_type(clk_type),
                                            ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    There are functions to properly convert between all of these types, use
    them so there are no longer any warnings.

    Fixes: a43913ea50a5 ("drm/amd/powerplay: add function get_clock_by_type_with_latency for
navi10")
    Fixes: e5e4e22391c2 ("drm/amd/powerplay: add interface to get clock by type with latency
for display (v2)")
    Link: https://github.com/ClangBuiltLinux/linux/issues/586
```

```
    Signed-off-by: Nathan Chancellor <natechancellor@gmail.com>
    Reviewed-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4326d0269e29eaf678d9079e2b2e58a587c4595e
Author: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
Date:   Mon Jul 15 18:04:08 2019 -0400

    drm/amdgpu: Fix hard hang for S/G display BOs.

    [ Upstream commit e4c4073b0139d055d43a9568690fc560aab4fa5c ]

    HW requires for caching to be unset for scanout BO
    mappings when the BO placement is in GTT memory.
    Usually the flag to unset is passed from user mode
    but for FB mode this was missing.

    v2:
    Keep all BO placement logic in amdgpu_display_supported_domains

    Suggested-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Tested-by: Shirish S <shirish.s@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f1835cadc8453caea6b06db2993b59ba008dba8
Author: Sean Paul <seanpaul@chromium.org>
Date:   Tue Jun 11 12:08:21 2019 -0400

    drm/rockchip: Check for fast link training before enabling psr

    [ Upstream commit ad309284a52be47c8b3126c9376358bf381861bc ]

    Once we start shutting off the link during PSR, we're going to want fast
    training to work. If the display doesn't support fast training, don't
    enable psr.

    Changes in v2:
    - None
    Changes in v3:
    - None
    Changes in v4:
    - None
    Changes in v5:
    - None

    Link to v1: https://patchwork.freedesktop.org/patch/msgid/20190228210939.83386-3-
sean@poorly.run
    Link to v2: https://patchwork.freedesktop.org/patch/msgid/20190326204509.96515-2-
sean@poorly.run
    Link to v3: https://patchwork.freedesktop.org/patch/msgid/20190502194956.218441-9-
sean@poorly.run
    Link to v4: https://patchwork.freedesktop.org/patch/msgid/20190508160920.144739-8-
sean@poorly.run

    Cc: Zain Wang <wzz@rock-chips.com>
    Cc: Tomasz Figa <tfiga@chromium.org>
    Tested-by: Heiko Stuebner <heiko@sntech.de>
    Reviewed-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sean Paul <seanpaul@chromium.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190611160844.257498-8-
sean@poorly.run
```

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

commit 22bae50f75344994bec3337795d52592921ffd08
Author: Navid Emamdoost <navid.emamdoost@gmail.com>
Date:   Wed Jul 24 14:55:34 2019 -0500

    drm/panel: check failure cases in the probe func

    [ Upstream commit afd6d4f5a52c16e1483328ac074abb1cde92c29f ]

    The following function calls may fail and return NULL, so the null check
    is added.
    of_graph_get_next_endpoint
    of_graph_get_remote_port_parent
    of_graph_get_remote_port

    Update: Thanks to Sam Ravnborg, for suggession on the use of goto to avoid
    leaking endpoint.

    Signed-off-by: Navid Emamdoost <navid.emamdoost@gmail.com>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190724195534.9303-1-
navid.emamdoost@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4e4e0e4954ed23ff8930c5b01c37e84769d00134
Author: Olivier Moysan <olivier.moysan@st.com>
Date:   Mon Jul 22 18:05:58 2019 +0200

    drm/bridge: sii902x: fix missing reference to mclk clock

    [ Upstream commit 365d28c92f8cd3d3860f8dd057a8c136e24b3698 ]

    Add devm_clk_get call to retrieve reference to master clock.

    Fixes: ff5781634c41 ("drm/bridge: sii902x: Implement HDMI audio support")

    Signed-off-by: Olivier Moysan <olivier.moysan@st.com>
    Reviewed-by: Jyri Sarha <jsarha@ti.com>
    Acked-by: Andrzej Hajda <a.hajda@samsung.com
    Signed-off-by: Benjamin Gaignard <benjamin.gaignard@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/1563811560-29589-2-git-send-email-
olivier.moysan@st.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 835bb6a305683387f4e8f878009f7d05ec12a9de
Author: Ahmad Fatoum <a.fatoum@pengutronix.de>
Date:   Fri Jul 12 10:42:28 2019 +0200

    drm/stm: attach gem fence to atomic state

    [ Upstream commit 8fabc9c3109a71b3577959a05408153ae69ccd8d ]

    To properly synchronize with other devices the fence from the GEM
    object backing the framebuffer needs to be attached to the atomic
    state, so the commit work can wait on fence signaling.

    Signed-off-by: Ahmad Fatoum <a.fatoum@pengutronix.de>
    Signed-off-by: Lucas Stach <l.stach@pengutronix.de>
    Acked-by: Philippe Cornu <philippe.cornu@st.com>
    Tested-by: Philippe Cornu <philippe.cornu@st.com>
    Signed-off-by: Benjamin Gaignard <benjamin.gaignard@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190712084228.8338-1-
l.stach@pengutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit c68a7f28126f8af7f477fa9ee039c5955f99ec77
Author: Noralf Trønnes <noralf@tronnes.org>
Date:   Mon Jul 22 12:43:09 2019 +0200

    drm/tinydrm/Kconfig: drivers: Select BACKLIGHT_CLASS_DEVICE

    [ Upstream commit 3389669ac5ea598562673c04971d7bb0fab0e9f1 ]

    The mipi_dbi helper is missing a dependency on DRM_KMS_HELPER and putting
    that in revealed this problem:

    drivers/video/fbdev/Kconfig:12:error: recursive dependency detected!
    drivers/video/fbdev/Kconfig:12: symbol FB is selected by DRM_KMS_FB_HELPER
    drivers/gpu/drm/Kconfig:75:     symbol DRM_KMS_FB_HELPER depends on DRM_KMS_HELPER
    drivers/gpu/drm/Kconfig:69:     symbol DRM_KMS_HELPER is selected by TINYDRM_MIPI_DBI
    drivers/gpu/drm/tinydrm/Kconfig:11:     symbol TINYDRM_MIPI_DBI is selected by
TINYDRM_HX8357D
    drivers/gpu/drm/tinydrm/Kconfig:15:     symbol TINYDRM_HX8357D depends on
BACKLIGHT_CLASS_DEVICE
    drivers/video/backlight/Kconfig:144:    symbol BACKLIGHT_CLASS_DEVICE is selected by
FB_BACKLIGHT
    drivers/video/fbdev/Kconfig:187:        symbol FB_BACKLIGHT depends on FB

    A symbol that selects DRM_KMS_HELPER can not depend on
    BACKLIGHT_CLASS_DEVICE. The reason for this is that DRM_KMS_FB_HELPER
    selects FB instead of depending on it.

    The tinydrm drivers have somehow gotten away with depending on
    BACKLIGHT_CLASS_DEVICE because DRM_TINYDRM selects DRM_KMS_HELPER and the
    drivers depend on that symbol.

    An audit shows that all DRM drivers that select DRM_KMS_HELPER and use
    BACKLIGHT_CLASS_DEVICE, selects it:
      DRM_TILCDC, DRM_GMA500, DRM_SHMOBILE, DRM_NOUVEAU, DRM_FSL_DCU,
      DRM_I915, DRM_RADEON, DRM_AMDGPU, DRM_PARADE_PS8622

    Documentation/kbuild/kconfig-language.txt has a note regarding select:
    1. 'select should be used with care since it doesn't visit dependencies.'
       This is not a problem since BACKLIGHT_CLASS_DEVICE doesn't have any
       dependencies.
    2. 'In general use select only for non-visible symbols'
       BACKLIGHT_CLASS_DEVICE is user visible.

    The real solution to this would be to have DRM_KMS_FB_HELPER depend on the
    user visible symbol FB. That is a can of worms I'm not willing to tackle.
    I fear that such a change will result in me handling difficult fallouts
    for the next weeks. So I'm following DRM suite here.

    Signed-off-by: Noralf Trønnes <noralf@tronnes.org>
    Reviewed-by: David Lechner <david@lechnology.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190722104312.16184-7-
noralf@tronnes.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35acba5957c9f715354774d44d4c06ab97594bb8
Author: Marko Kohtala <marko.kohtala@okoko.fi>
Date:   Tue Jun 18 10:41:08 2019 +0300

    video: ssd1307fb: Start page range at page_offset

    [ Upstream commit dd9782834dd9dde3624ff1acea8859f3d3e792d4 ]

    The page_offset was only applied to the end of the page range. This caused
    the display updates to cause a scrolling effect on the display because the
```

```
    amount of data written to the display did not match the range display
    expected.

    Fixes: 301bc0675b67 ("video: ssd1307fb: Make use of horizontal addressing mode")
    Signed-off-by: Marko Kohtala <marko.kohtala@okoko.fi>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Rob Herring <robh+dt@kernel.org>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Cc: David Airlie <airlied@linux.ie>
    Cc: Michal Vokáč <michal.vokac@ysoft.com>
    Signed-off-by: Bartlomiej Zolnierkiewicz <b.zolnierkie@samsung.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190618074111.9309-4-
marko.kohtala@okoko.fi
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1eb67506dd3376a167519b8fafd5a85c79fe3cba
Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
Date:   Thu Jul 11 12:32:43 2019 -0400

    drm/amd/display: Copy GSL groups when committing a new context

    [ Upstream commit 21ffcc94d5b3dc024fedac700f1e7f9dacf4ab4f ]

    [Why]
    DC configures the GSL group for the pipe when pipe_split is enabled
    and we're switching flip types (buffered <-> immediate flip) on DCN2.

    In order to record what GSL group the pipe is using DC stores it in
    the pipe's stream_res. DM is not aware of this internal grouping, nor
    is DC resource.

    So when DM creates a dc_state context and passes it to DC the current
    GSL group is lost - DM never knew about it in the first place.

    After 3 immediate flips we run out of GSL groups and we're no longer
    able to correctly perform *any* flip for multi-pipe scenarios.

    [How]
    The gsl_group needs to be copied to the new context.

    DM has no insight into GSL grouping and could even potentially create
    a brand new context without referencing current hardware state. So this
    makes the most sense to have happen in DC.

    There are two places where DC can apply a new context:
    - dc_commit_state
    - dc_commit_updates_for_stream

    But what's shared between both of these is apply_ctx_for_surface.

    This logic only matters for DCN2, so it can be placed in
    dcn20_apply_ctx_for_surface. Before doing any locking (where the GSL
    group is setup) we can copy over the GSL groups before committing the
    new context.

    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Reviewed-by: Hersen Wu <hersen.wu@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 67cdde3f8588723932b510d41313a70482cdfa2a
Author: Nikola Cornij <nikola.cornij@amd.com>
Date:   Tue Jun 25 17:19:25 2019 -0400
```

```
    drm/amd/display: Clear FEC_READY shadow register if DPCD write fails

    [ Upstream commit d68a74541735e030dea56f72746cd26d19986f41 ]

    [why]
    As a fail-safe, in case 'set FEC_READY' DPCD write fails, a HW shadow
    register should be cleared and the internal FEC stat should be set to
    'not ready'. This is to make sure HW settings will be consistent with
    FEC_READY state on the RX.

    Signed-off-by: Nikola Cornij <nikola.cornij@amd.com>
    Reviewed-by: Joshua Aberback <Joshua.Aberback@amd.com>
    Acked-by: Chris Park <Chris.Park@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a347635c86e94e9e90780d060acd587b1454321a
Author: Su Sung Chung <Su.Chung@amd.com>
Date:   Fri Jun 21 16:14:36 2019 -0400

    drm/amd/display: fix not calling ppsmu to trigger PME

    [ Upstream commit 18b401874aee10c80b5745c9b93280dae5a59809 ]

    [why]
    dcn20_clk_mgr_construct was not initializing pp_smu, and PME call gets
    filtered out by the null check

    [how]
    initialize pp_smu dcn20_clk_mgr_construct

    Signed-off-by: Su Sung Chung <Su.Chung@amd.com>
    Reviewed-by: Eric Yang <eric.yang2@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90d92ba20d71df34b8a55b8e71d20cba608957c8
Author: Nikola Cornij <nikola.cornij@amd.com>
Date:   Wed Jun 19 14:30:52 2019 -0400

    drm/amd/display: Power-gate all DSCs at driver init time

    [ Upstream commit 75c35000235f3662f2810e9a59b0c8eed045432e ]

    [why]
    DSC should be powered-on only on as-needed basis, i.e. if the mode
    requires it

    [how]
    Loop over all the DSCs at driver init time and power-gate each

    Signed-off-by: Nikola Cornij <nikola.cornij@amd.com>
    Reviewed-by: Nevenko Stupar <Nevenko.Stupar@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 13a5e5e123c7fa14af5703dcceda9ab7e9c9e236
Author: Anthony Koo <anthony.koo@amd.com>
Date:   Wed Jun 12 16:30:15 2019 -0400

    drm/amd/display: add monitor patch to add T7 delay
```

```
    [ Upstream commit 88eac241a1fc500ce5274a09ddc4bd5fc2b5adb6 ]

    [Why]
    Specifically to one panel,
    TCON is able to accept active video signal quickly, but
    the Source Driver requires 2-3 frames of extra time.

    It is a Panel issue since TCON needs to take care of
    all Sink requirements including Source Driver. But in
    this case it does not.

    Customer is asking to add fixed T7 delay as panel
    workaround.

    [How]
    Add monitor specific patch to add T7 delay

    Signed-off-by: Anthony Koo <anthony.koo@amd.com>
    Reviewed-by: Charlene Liu <Charlene.Liu@amd.com>
    Acked-by: Leo Li <sunpeng.li@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2dbb9e65378c6e8be864f1ebfe0aaa3d0b262b58
Author: Lucas Stach <l.stach@pengutronix.de>
Date:   Wed Jul 10 15:07:40 2019 +0200

    drm/panel: simple: fix AUO g185han01 horizontal blanking

    [ Upstream commit f8c6bfc612b56f02e1b8fae699dff12738aaf889 ]

    The horizontal blanking periods are too short, as the values are
    specified for a single LVDS channel. Since this panel is dual LVDS
    they need to be doubled. With this change the panel reaches its
    nominal vrefresh rate of 60Fps, instead of the 64Fps with the
    current wrong blanking.

    Philipp Zabel added:
    The datasheet specifies 960 active clocks + 40/128/160 clocks blanking
    on each of the two LVDS channels (min/typical/max), so doubled this is
    now correct.

    Signed-off-by: Lucas Stach <l.stach@pengutronix.de>
    Reviewed-by: Philipp Zabel <p.zabel@pengutronix.de>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/1562764060.23869.12.camel@pengutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2e0be4a3f74e30982ce445ca2f66acbbb10e9fa9
Author: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
Date:   Tue Jun 25 22:36:18 2019 -0300

    drm/vkms: Avoid assigning 0 for possible_crtc

    [ Upstream commit e9d85f731de06a35d2ae6cdcf7d0e037c98ef41a ]

    When vkms invoke drm_universal_plane_init(), it sets 0 for
    possible_crtcs parameter which means that planes can't be attached to
    any CRTC. It currently works due to some safeguard in the drm_crtc file;
    however, it is possible to identify the problem by trying to append a
    second connector. This patch fixes this issue by modifying
    vkms_plane_init() to accept an index parameter which makes the code a
    little bit more flexible and avoid set zero to possible_crtcs.
```

```
    Signed-off-by: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
    Reviewed-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link:
https://patchwork.freedesktop.org/patch/msgid/d67849c62a8d8ace1a0af455998b588798a4c45f.15614
91964.git.rodrigosiqueiramelo@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19feb86ed0d5a39a0d1a450707b2f63137d5fa5f
Author: Andrey Smirnov <andrew.smirnov@gmail.com>
Date:   Tue Jun 18 22:27:09 2019 -0700

    drm/bridge: tc358767: Increase AUX transfer length limit

    [ Upstream commit e0655feaec62d5139b6b13a7b1bbb1ab8f1c2d83 ]

    According to the datasheet tc358767 can transfer up to 16 bytes via
    its AUX channel, so the artificial limit of 8 appears to be too
    low. However only up to 15-bytes seem to be actually supported and
    trying to use 16-byte transfers results in transfers failing
    sporadically (with bogus status in case of I2C transfers), so limit it
    to 15.

    Signed-off-by: Andrey Smirnov <andrew.smirnov@gmail.com>
    Reviewed-by: Andrzej Hajda <a.hajda@samsung.com>
    Reviewed-by: Tomi Valkeinen <tomi.valkeinen@ti.com>
    Cc: Andrzej Hajda <a.hajda@samsung.com>
    Cc: Laurent Pinchart <Laurent.pinchart@ideasonboard.com>
    Cc: Tomi Valkeinen <tomi.valkeinen@ti.com>
    Cc: Andrey Gusakov <andrey.gusakov@cogentembedded.com>
    Cc: Philipp Zabel <p.zabel@pengutronix.de>
    Cc: Cory Tusar <cory.tusar@zii.aero>
    Cc: Chris Healy <cphealy@gmail.com>
    Cc: Lucas Stach <l.stach@pengutronix.de>
    Cc: dri-devel@lists.freedesktop.org
    Cc: linux-kernel@vger.kernel.org
    Signed-off-by: Andrzej Hajda <a.hajda@samsung.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190619052716.16831-9-
andrew.smirnov@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6ee696a950e7a0afd7dbd42483e85e7307951fbc
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Tue Jun 18 13:52:45 2019 +0200

    drm/mcde: Fix uninitialized variable

    [ Upstream commit ca5be902a87ddccc88144f2dea21b5f0814391ef ]

    We need to handle the case when of_drm_find_bridge() returns
    NULL.

    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Acked-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190618115245.13915-1-
linus.walleij@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4abf9e52d7a40c32bb88f17e7f79adc2bc3b713b
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Fri Jun 7 00:27:42 2019 +0200

    drm/vkms: Fix crc worker races
```

```
    [ Upstream commit 18d0952a838ba559655b0cd9cf85097ad63d9bca ]

    The issue we have is that the crc worker might fall behind. We've
    tried to handle this by tracking both the earliest frame for which it
    still needs to compute a crc, and the last one. Plus when the
    crtc_state changes, we have a new work item, which are all run in
    order due to the ordered workqueue we allocate for each vkms crtc.

    Trouble is there's been a few small issues in the current code:
    - we need to capture frame_end in the vblank hrtimer, not in the
      worker. The worker might run much later, and then we generate a lot
      of crc for which there's already a different worker queued up.
    - frame number might be 0, so create a new crc_pending boolean to
      track this without confusion.
    - we need to atomically grab frame_start/end and clear it, so do that
      all in one go. This is not going to create a new race, because if we
      race with the hrtimer then our work will be re-run.
    - only race that can happen is the following:
      1. worker starts
      2. hrtimer runs and updates frame_end
      3. worker grabs frame_start/end, already reading the new frame_end,
      and clears crc_pending
      4. hrtimer calls queue_work()
      5. worker completes
      6. worker gets  re-run, crc_pending is false
      Explain this case a bit better by rewording the comment.

    v2: Demote warning level output to debug when we fail to requeue, this
    is expected under high load when the crc worker can't quite keep up.

    Cc: Shayenne Moura <shayenneluzmoura@gmail.com>
    Cc: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
    Cc: Haneen Mohammed <hamohammed.sa@gmail.com>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Reviewed-by: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
    Tested-by: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
    Signed-off-by: Rodrigo Siqueira <rodrigosiqueiramelo@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20190606222751.32567-2-
daniel.vetter@ffwll.ch
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```