```
commit 313c8460cf0290fb1b9f71a20573fc32ac6c9cee
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Sat Feb 1 09:34:53 2020 +0000

    Linux 5.4.17

commit 831919b4c481c4854fd75d87190b0c42e769cf35
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sat Nov 16 14:56:19 2019 +0100

    power/supply: ingenic-battery: Don't change scale if there's only one

    commit 86b9182df8bb12610d4d6feac45a69f3ed57bfd2 upstream.

    The ADC in the JZ4740 can work either in high-precision mode with a 2.5V
    range, or in low-precision mode with a 7.5V range. The code in place in
    this driver will select the proper scale according to the maximum
    voltage of the battery.

    The JZ4770 however only has one mode, with a 6.6V range. If only one
    scale is available, there's no need to change it (and nothing to change
    it to), and trying to do so will fail with -EINVAL.

    Fixes: fb24ccfbe1e0 ("power: supply: add Ingenic JZ47xx battery driver.")

    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Acked-by: Artur Rojek <contact@artur-rojek.eu>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dd350f3918be4b2bb298647cd9d467beeaf2c22a
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Dec 4 17:43:46 2019 +0100

    Revert "um: Enable CONFIG_CONSTRUCTORS"

    commit 87c9366e17259040a9118e06b6dc8de986e5d3d1 upstream.

    This reverts commit 786b2384bf1c ("um: Enable CONFIG_CONSTRUCTORS").

    There are two issues with this commit, uncovered by Anton in tests
    on some (Debian) systems:

    1) I completely forgot to call any constructors if CONFIG_CONSTRUCTORS
       isn't set. Don't recall now if it just wasn't needed on my system, or
       if I never tested this case.

    2) With that fixed, it works - with CONFIG_CONSTRUCTORS *unset*. If I
       set CONFIG_CONSTRUCTORS, it fails again, which isn't totally
       unexpected since whatever wanted to run is likely to have to run
       before the kernel init etc. that calls the constructors in this case.

    Basically, some constructors that gcc emits (libc has?) need to run
    very early during init; the failure mode otherwise was that the ptrace
    fork test already failed:

    ----------------------
    $ ./linux mem=512M
    Core dump limits :
            soft - 0
            hard - NONE
    Checking that ptrace can change system call numbers...check_ptrace : child exited with
exitcode 6, while expecting 0; status 0x67f
```

```
    Aborted
    ----------------------

    Thinking more about this, it's clear that we simply cannot support
    CONFIG_CONSTRUCTORS in UML. All the cases we need now (gcov, kasan)
    involve not use of the __attribute__((constructor)), but instead
    some constructor code/entry generated by gcc. Therefore, we cannot
    distinguish between kernel constructors and system constructors.

    Thus, revert this commit.

    Cc: stable@vger.kernel.org [5.4+]
    Fixes: 786b2384bf1c ("um: Enable CONFIG_CONSTRUCTORS")
    Reported-by: Anton Ivanov <anton.ivanov@cambridgegreys.com>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Acked-by: Anton Ivanov <anton.ivanov@cambridgegreys.co.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

    Signed-off-by: Richard Weinberger <richard@nod.at>

commit a775f67b4b6110d07bea24fcd3303ca53fcb0ecf
Author: Andrew Murray <andrew.murray@arm.com>
Date:   Fri Jan 17 13:43:24 2020 +0000

    KVM: arm64: Write arch.mdcr_el2 changes since last vcpu_load on VHE

    commit 4942dc6638b07b5326b6d2faa142635c559e7cd5 upstream.

    On VHE systems arch.mdcr_el2 is written to mdcr_el2 at vcpu_load time to
    set options for self-hosted debug and the performance monitors
    extension.

    Unfortunately the value of arch.mdcr_el2 is not calculated until
    kvm_arm_setup_debug() in the run loop after the vcpu has been loaded.
    This means that the initial brief iterations of the run loop use a zero
    value of mdcr_el2 - until the vcpu is preempted. This also results in a
    delay between changes to vcpu->guest_debug taking effect.

    Fix this by writing to mdcr_el2 in kvm_arm_setup_debug() on VHE systems
    when a change to arch.mdcr_el2 has been detected.

    Fixes: d5a21bcc2995 ("KVM: arm64: Move common VHE/non-VHE trap config in separate
functions")
    Cc: <stable@vger.kernel.org> # 4.17.x-
    Suggested-by: James Morse <james.morse@arm.com>
    Acked-by: Will Deacon <will@kernel.org>
    Reviewed-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Andrew Murray <andrew.murray@arm.com>
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 00d58d0b58bc721990215874f1e1e1a491846ebb
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Tue Nov 19 17:41:31 2019 +0800

    crypto: pcrypt - Fix user-after-free on module unload

    commit 07bfd9bdf568a38d9440c607b72342036011f727 upstream.

    On module unload of pcrypt we must unregister the crypto algorithms
    first and then tear down the padata structure.  As otherwise the
    crypto algorithms are still alive and can be used while the padata
    structure is being freed.

    Fixes: 5068c7a883d1 ("crypto: pcrypt - Add pcrypt crypto...")
```

```
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f6d555f58f3fa3c610ec08b9bd5d6af7ad0592d
Author: Iuliana Prodan <iuliana.prodan@nxp.com>
Date:   Wed Nov 27 00:54:26 2019 +0200

    crypto: caam - do not reset pointer size from MCFGR register

    commit 7278fa25aa0ebcc0e62c39b12071069df13f7e77 upstream.

    In commit 'a1cf573ee95 ("crypto: caam - select DMA address size at runtime")'
    CAAM pointer size (caam_ptr_size) is changed from
    sizeof(dma_addr_t) to runtime value computed from MCFGR register.
    Therefore, do not reset MCFGR[PS].

    Fixes: a1cf573ee95 ("crypto: caam - select DMA address size at runtime")
    Signed-off-by: Iuliana Prodan <iuliana.prodan@nxp.com>
    Cc: <stable@vger.kernel.org>
    Cc: Andrey Smirnov <andrew.smirnov@gmail.com>
    Cc: Alison Wang <alison.wang@nxp.com>
    Reviewed-by: Horia Geantă <horia.geanta@nxp.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b52ada5b82cb7cb83190482b9383509c998ba8b
Author: Daniel Axtens <dja@axtens.net>
Date:   Wed Jan 8 16:06:46 2020 +1100

    crypto: vmx - reject xts inputs that are too short

    commit 1372a51b88fa0d5a8ed2803e4975c98da3f08463 upstream.

    When the kernel XTS implementation was extended to deal with ciphertext
    stealing in commit 8083b1bf8163 ("crypto: xts - add support for ciphertext
    stealing"), a check was added to reject inputs that were too short.

    However, in the vmx enablement - commit 239668419349 ("crypto: vmx/xts -
    use fallback for ciphertext stealing"), that check wasn't added to the
    vmx implementation. This disparity leads to errors like the following:

    alg: skcipher: p8_aes_xts encryption unexpectedly succeeded on test vector "random:
len=0 klen=64"; expected_error=-22, cfg="random: inplace may_sleep use_finup src_divs=
[<flush>66.99%@+10, 33.1%@alignmask+1155]"

    Return -EINVAL if asked to operate with a cryptlen smaller than the AES
    block size. This brings vmx in line with the generic implementation.

    Reported-by: Erhard Furtner <erhard_f@mailbox.org>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=206049
    Fixes: 239668419349 ("crypto: vmx/xts - use fallback for ciphertext stealing")
    Cc: Ard Biesheuvel <ard.biesheuvel@linaro.org>
    Cc: stable@vger.kernel.org # v5.4+
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    [dja: commit message]
    Signed-off-by: Daniel Axtens <dja@axtens.net>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9d9b578ffef37278d8c37492c6a0fe75e1f6bab8
Author: Herbert Xu <herbert@gondor.apana.org.au>
Date:   Thu Dec 5 13:45:05 2019 +0800

    crypto: af_alg - Use bh_lock_sock in sk_destruct
```

```
    commit 37f96694cf73ba116993a9d2d99ad6a75fa7fdb0 upstream.

    As af_alg_release_parent may be called from BH context (most notably
    due to an async request that only completes after socket closure,
    or as reported here because of an RCU-delayed sk_destruct call), we
    must use bh_lock_sock instead of lock_sock.

    Reported-by: syzbot+c2f1558d49e25cc36e5e@syzkaller.appspotmail.com
    Reported-by: Eric Dumazet <eric.dumazet@gmail.com>
    Fixes: c840ac6af3f8 ("crypto: af_alg - Disallow bind/setkey/...")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ac39443b7b497f79891b16edb0f468ba35d36aee
Author: Johan Hovold <johan@kernel.org>
Date:   Thu Nov 28 18:22:03 2019 +0100

    rsi: fix non-atomic allocation in completion handler

    commit b9b9f9fea21830f85cf0148cd8dce001ae55ead1 upstream.

    USB completion handlers are called in atomic context and must
    specifically not allocate memory using GFP_KERNEL.

    Fixes: a1854fae1414 ("rsi: improve RX packet handling in USB interface")
    Cc: stable <stable@vger.kernel.org> # 4.17
    Cc: Prameela Rani Garnepudi <prameela.j04cs@gmail.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1f04209d19c505be7e4fd6b39c8b776819e80555
Author: Johan Hovold <johan@kernel.org>
Date:   Thu Nov 28 18:22:02 2019 +0100

    rsi: fix memory leak on failed URB submission

    commit 47768297481184932844ab01a86752ba31a38861 upstream.

    Make sure to free the skb on failed receive-URB submission (e.g. on
    disconnect or currently also due to a missing endpoint).

    Fixes: a1854fae1414 ("rsi: improve RX packet handling in USB interface")
    Cc: stable <stable@vger.kernel.org>     # 4.17
    Cc: Prameela Rani Garnepudi <prameela.j04cs@gmail.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2d367468ecdf3fb58e58bfb534cc8c97d1b2eae5
Author: Johan Hovold <johan@kernel.org>
Date:   Thu Nov 28 18:22:01 2019 +0100

    rsi: fix use-after-free on probe errors

    commit 92aafe77123ab478e5f5095878856ab0424910da upstream.

    The driver would fail to stop the command timer in most error paths,
    something which specifically could lead to the timer being freed while
    still active on I/O errors during probe.

    Fix this by making sure that each function starting the timer also stops
    it in all relevant error paths.
```

```
    Reported-by: syzbot+1d1597a5aa3679c65b9f@syzkaller.appspotmail.com
    Fixes: b78e91bcfb33 ("rsi: Add new firmware loading method")
    Cc: stable <stable@vger.kernel.org>     # 4.12
    Cc: Prameela Rani Garnepudi <prameela.j04cs@gmail.com>
    Cc: Amitkumar Karwar <amit.karwar@redpinesignals.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d8902b43745fa2d71cddda29f079f40ce1b76a11
Author: Johan Hovold <johan@kernel.org>
Date:   Thu Nov 28 18:22:00 2019 +0100

    rsi: fix use-after-free on failed probe and unbind

    commit e93cd35101b61e4c79149be2cfc927c4b28dc60c upstream.

    Make sure to stop both URBs before returning after failed probe as well
    as on disconnect to avoid use-after-free in the completion handler.

    Reported-by: syzbot+b563b7f8dbe8223a51e8@syzkaller.appspotmail.com
    Fixes: a4302bff28e2 ("rsi: add bluetooth rx endpoint")
    Fixes: dad0d04fa7ba ("rsi: Add RS9113 wireless driver")
    Cc: stable <stable@vger.kernel.org>     # 3.15
    Cc: Siva Rebbagondla <siva.rebbagondla@redpinesignals.com>
    Cc: Prameela Rani Garnepudi <prameela.j04cs@gmail.com>
    Cc: Amitkumar Karwar <amit.karwar@redpinesignals.com>
    Cc: Fariya Fatima <fariyaf@gmail.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 99d9138295cbbbb241ae1997aa847a44d0a65c48
Author: Tony Lindgren <tony@atomide.com>
Date:   Mon Dec 9 13:55:15 2019 -0800

    bus: ti-sysc: Fix missing force mstandby quirk handling

    [ Upstream commit 93c60483b5feefced92b869d5f97769495bc6313 ]

    Commit 03856e928b0e ("bus: ti-sysc: Handle mstandby quirk and use it for
    musb") added quirk handling for mstandby quirk but did not consider that
    we also need a quirk variant for SYSC_QUIRK_FORCE_MSTANDBY.

    We need to use forced idle mode for both SYSC_QUIRK_SWSUP_MSTANDBY and
    SYSC_QUIRK_FORCE_MSTANDBY, but SYSC_QUIRK_SWSUP_MSTANDBY also need to
    additionally also configure no-idle mode when enabled.

    Fixes: 03856e928b0e ("bus: ti-sysc: Handle mstandby quirk and use it for musb")
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58b09300ee1b505d125da7f0d8f0f8641627e630
Author: Andre Heider <a.heider@gmail.com>
Date:   Fri Nov 22 13:31:42 2019 +0100

    Bluetooth: btbcm: Use the BDADDR_PROPERTY quirk

    [ Upstream commit a4f95f31a9f38d9bb1fd313fcc2d0c0d48116ee3 ]

    Some devices ship with the controller default address, like the
    Orange Pi 3 (BCM4345C5).

    Allow the bootloader to set a valid address through the device tree.
```

```
    Signed-off-by: Andre Heider <a.heider@gmail.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1e5f8c5156e1f957cf3e3b065836566467489d1
Author: Marcel Holtmann <marcel@holtmann.org>
Date:   Fri Nov 22 00:33:45 2019 +0100

    Bluetooth: Allow combination of BDADDR_PROPERTY and INVALID_BDADDR quirks

    [ Upstream commit 7fdf6c6a0d0e032aac2aa4537a23af1e04a397ce ]

    When utilizing BDADDR_PROPERTY and INVALID_BDADDR quirks together it
    results in an unconfigured controller even if the bootloader provides
    a valid address. Fix this by allowing a bootloader provided address
    to mark the controller as configured.

    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Tested-by: Andre Heider <a.heider@gmail.com>
    Signed-off-by: Johan Hedberg <johan.hedberg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d273ee188a4a8fa558cdfdaad9dfaa3968259e21
Author: Hui Wang <hui.wang@canonical.com>
Date:   Thu Nov 21 10:26:44 2019 +0800

    ALSA: hda/realtek - Move some alc236 pintbls to fallback table

    [ Upstream commit d64ebdbfd4f71406f58210f5ccb16977b4cd31d2 ]

    We have a new Dell machine which needs to apply the quirk
    ALC255_FIXUP_DELL1_MIC_NO_PRESENCE, try to use the fallback table
    to fix it this time. And we could remove all pintbls of alc236
    for applying DELL1_MIC_NO_PRESENCE on Dell machines.

    Signed-off-by: Hui Wang <hui.wang@canonical.com>
    Link: https://lore.kernel.org/r/20191121022644.8078-2-hui.wang@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 050072c0f7988beb75afb9d4c6057618e64617e4
Author: Laura Abbott <labbott@fedoraproject.org>
Date:   Tue Sep 8 09:53:38 2015 -0700

    usb-storage: Disable UAS on JMicron SATA enclosure

    [ Upstream commit bc3bdb12bbb3492067c8719011576370e959a2e6 ]

    Steve Ellis reported incorrect block sizes and alignement
    offsets with a SATA enclosure. Adding a quirk to disable
    UAS fixes the problems.

    Reported-by: Steven Ellis <sellis@redhat.com>
    Cc: Pacho Ramos <pachoramos@gmail.com>
    Signed-off-by: Laura Abbott <labbott@fedoraproject.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a2e6217bd8ea5d5f265446a7883f972df7ee0294
Author: Tony Lindgren <tony@atomide.com>
Date:   Thu Nov 14 08:25:28 2019 -0800

    bus: ti-sysc: Add module enable quirk for audio AESS
```

```
    [ Upstream commit 020003f763e24e4ed0bb3d8909f3940891536d5d ]

    We must set the autogating bit on enable for AESS (Audio Engine SubSystem)
    when probed with ti-sysc interconnect target module driver. Otherwise it
    won't idle properly.

    Cc: Peter Ujfalusi <peter.ujfalusi@ti.com>
    Tested-by: Peter Ujfalusi <peter.ujfalusi@ti.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c101a22fec5c612fc5aa9d91c8cf14b477a5cda
Author: Adrian Hunter <adrian.hunter@intel.com>
Date:   Thu Oct 10 15:46:30 2019 +0300

    mmc: sdhci-pci: Add support for Intel JSL

    [ Upstream commit 315e3bd7ac19b18ba704d96cbb9b79bad485c01f ]

    Add PCI Ids for Intel JSL.

    Signed-off-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac342e0a1f186954013024079da24b5beaa9276d
Author: Raul E Rangel <rrangel@chromium.org>
Date:   Wed Sep 4 10:46:25 2019 -0600

    mmc: sdhci-pci: Quirk for AMD SDHC Device 0x7906

    [ Upstream commit 7a869f00bb15bcefb8804d798a49b086267b03e6 ]

    AMD SDHC 0x7906 requires a hard reset to clear all internal state.
    Otherwise it can get into a bad state where the DATA lines are always
    read as zeros.

    This change requires firmware that can transition the device into
    D3Cold for it to work correctly. If the firmware does not support
    transitioning to D3Cold then the power state transitions are a no-op.

    Signed-off-by: Raul E Rangel <rrangel@chromium.org>
    Signed-off-by: Shyam Sundar S K <Shyam-sundar.S-k@amd.com>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 142356a63872028eb46ca9154ceaba4eace3418c
Author: Ben Dooks <ben.dooks@codethink.co.uk>
Date:   Tue Oct 15 14:12:38 2019 +0100

    ARM: OMAP2+: SmartReflex: add omap_sr_pdata definition

    [ Upstream commit 2079fe6ea8cbd2fb2fbadba911f1eca6c362eb9b ]

    The omap_sr_pdata is not declared but is exported, so add a
    define for it to fix the following warning:

    arch/arm/mach-omap2/pdata-quirks.c:609:36: warning: symbol 'omap_sr_pdata' was not
declared. Should it be static?

    Signed-off-by: Ben Dooks <ben.dooks@codethink.co.uk>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit f3020d789311df570a40e5112959b15597c20444
Author: Joel Stanley <joel@jms.id.au>
Date:   Tue Nov 12 10:39:26 2019 +1030

    ARM: config: aspeed-g5: Enable 8250_DW quirks

    [ Upstream commit a5331a7a87ec81d5228b7421acf831b2d0c0de26 ]

    This driver option is used by the AST2600 A0 boards to work around a
    hardware issue.

    Reviewed-by: Cédric Le Goater <clg@kaod.org>
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Joel Stanley <joel@jms.id.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f8f8f31c50574ef2930b0042f95761188a052db
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Tue Oct 29 11:44:09 2019 +0200

    mfd: intel-lpss: Add Intel Comet Lake PCH-H PCI IDs

    [ Upstream commit dd047dce3a6f5233b98e792e2287cc549da35879 ]

    Intel Comet Lake PCH-H has the same LPSS than Intel Cannon Lake.
    Add the new IDs to the list of supported devices.

    Signed-off-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 81909bd070db431e9ef271aa58d2d6fb9df91616
Author: Joakim Zhang <qiangqing.zhang@nxp.com>
Date:   Fri Nov 1 08:36:16 2019 +0000

    perf/imx_ddr: Add enhanced AXI ID filter support

    [ Upstream commit 44f8bd014a94ed679ddb77d0b92350d4ac4f23a5 ]

    With DDR_CAP_AXI_ID_FILTER quirk, indicating HW supports AXI ID filter
    which only can get bursts from DDR transaction, i.e. DDR read/write
    requests.

    This patch add DDR_CAP_AXI_ID_ENHANCED_FILTER quirk, indicating HW
    supports AXI ID filter which can get bursts and bytes from DDR
    transaction at the same time. We hope PMU always return bytes in the
    driver due to it is more meaningful for users.

    Signed-off-by: Joakim Zhang <qiangqing.zhang@nxp.com>
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b02b0a6bcc81c76d243eddcfed7271c8e54fde40
Author: Logan Gunthorpe <logang@deltatee.com>
Date:   Tue Oct 22 16:01:21 2019 -0600

    iommu/amd: Support multiple PCI DMA aliases in IRQ Remapping

    [ Upstream commit 3c124435e8dd516df4b2fc983f4415386fd6edae ]

    Non-Transparent Bridge (NTB) devices (among others) may have many DMA
    aliases seeing the hardware will send requests with different device ids
    depending on their origin across the bridged hardware.

    See commit ad281ecf1c7d ("PCI: Add DMA alias quirk for Microsemi Switchtec
```

```
    NTB") for more information on this.

    The AMD IOMMU IRQ remapping functionality ignores all PCI aliases for
    IRQs so if devices send an interrupt from one of their aliases they
    will be blocked on AMD hardware with the IOMMU enabled.

    To fix this, ensure IRQ remapping is enabled for all aliases with
    MSI interrupts.

    This is analogous to the functionality added to the Intel IRQ remapping
    code in commit 3f0c625c6ae7 ("iommu/vt-d: Allow interrupts from the entire
    bus for aliased devices")

    Signed-off-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f03a258f20f1699ede29bb40804074db9398a0d
Author: Logan Gunthorpe <logang@deltatee.com>
Date:   Tue Oct 22 16:01:20 2019 -0600

    iommu/amd: Support multiple PCI DMA aliases in device table

    [ Upstream commit 3332364e4ebc0581d133a334645a20fd13b580f1 ]

    Non-Transparent Bridge (NTB) devices (among others) may have many DMA
    aliases seeing the hardware will send requests with different device ids
    depending on their origin across the bridged hardware.

    See commit ad281ecf1c7d ("PCI: Add DMA alias quirk for Microsemi
    Switchtec NTB") for more information on this.

    The AMD IOMMU ignores all the PCI aliases except the last one so DMA
    transfers from these aliases will be blocked on AMD hardware with the
    IOMMU enabled.

    To fix this, ensure the DTEs are cloned for every PCI alias. This is
    done by copying the DTE data for each alias as well as the IVRS alias
    every time it is changed.

    Signed-off-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d7937ea02315cdd07348e064d71ea73e00b7ed03
Author: Jarkko Nikula <jarkko.nikula@linux.intel.com>
Date:   Tue Oct 29 13:58:02 2019 +0200

    spi: pxa2xx: Add support for Intel Comet Lake-H

    [ Upstream commit f0cf17ed76cffa365001d263ced1f130ec794917 ]

    Add Intel Comet Lake-H LPSS SPI PCI IDs.

    Signed-off-by: Jarkko Nikula <jarkko.nikula@linux.intel.com>
    Link: https://lore.kernel.org/r/20191029115802.6779-1-jarkko.nikula@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c90ed6e662881a43b455a00571d84d38b954b738
Author: Tony Lindgren <tony@atomide.com>
Date:   Mon Oct 21 14:15:55 2019 -0700

    bus: ti-sysc: Use swsup quirks also for am335x musb
```

```
    [ Upstream commit 1819ef2e2d12d5b1a6ee54ac1c2afe35cffc677c ]

    Also on am335x we need the swsup quirks for musb.

    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48b038c9d54a28cd49986a3dc8b6edf7fb21530c
Author: Tony Lindgren <tony@atomide.com>
Date:   Mon Oct 21 13:36:09 2019 -0700

    bus: ti-sysc: Handle mstandby quirk and use it for musb

    [ Upstream commit 03856e928b0e1a1c274eece1dfe4330a362c37f3 ]

    We need swsup quirks for sidle and mstandby for musb to work
    properly.

    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d34eacd8348cc8df9b261d5c6fd06d691b97784
Author: Thomas Voegtle <tv@lio96.de>
Date:   Wed Oct 2 14:26:03 2019 -0300

    media: dvbsky: add support for eyeTV Geniatech T2 lite

    [ Upstream commit 144945833368806406543 00c76d0f5df3360d85f ]

    Adds USB ID for the eyeTV Geniatech T2 lite to the dvbsky driver.
    This is a Geniatech T230C based stick without IR and a different USB ID.

    Signed-off-by: Thomas Voegtle <tv@lio96.de>
    Tested-by: Jan Pieter van Woerkom <jp@jpvw.nl>
    Signed-off-by: Sean Young <sean@mess.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+samsung@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fef41f3a82eaca01749478ddf857e8dfab6d0dc3
Author: Slawomir Pawlowski <slawomir.pawlowski@intel.com>
Date:   Tue Sep 17 09:20:48 2019 +0000

    PCI: Add DMA alias quirk for Intel VCA NTB

    [ Upstream commit 56b4cd4b7da9ee95778eb5c8abea49f641ebfd91 ]

    Intel Visual Compute Accelerator (VCA) is a family of PCIe add-in devices
    exposing computational units via Non Transparent Bridges (NTB, PEX 87xx).

    Similarly to MIC x200, we need to add DMA aliases to allow buffer access
    when IOMMU is enabled.

    Add aliases to allow computational unit access to host memory.  These
    aliases mark the whole VCA device as one IOMMU group.

    All possible slot numbers (0x20) are used, since we are unable to tell what
    slot is used on other side.  This quirk is intended for both host and
    computational unit sides.  The VCA devices have up to five functions: four
    for DMA channels and one additional.

    Link:
https://lore.kernel.org/r/5683A335CC8BE1438C3C30C49DCC38DF637CED8E@IRSMSX102.ger.corp.intel.
com
    Signed-off-by: Slawomir Pawlowski <slawomir.pawlowski@intel.com>
    Signed-off-by: Przemek Kitszel <przemyslawx.kitszel@intel.com>
```

```
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8aaefd5510ab7313e8f02d6c66e842cd7486825
Author: Pacien TRAN-GIRARD <pacien.trangirard@pacien.net>
Date:   Fri Sep 27 23:19:03 2019 +0200

    platform/x86: dell-laptop: disable kbd backlight on Inspiron 10xx

    [ Upstream commit 10b65e2915b2fcc606d173e98a972850101fb4c4 ]

    This patch adds a quirk disabling keyboard backlight support for the
    Dell Inspiron 1012 and 1018.

    Those models wrongly report supporting keyboard backlight control
    features (through SMBIOS tokens) even though they're not equipped with
    a backlit keyboard. This led to broken controls being exposed
    through sysfs by this driver which froze the system when used.

    Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=107651
    Signed-off-by: Pacien TRAN-GIRARD <pacien.trangirard@pacien.net>
    Reviewed-by: Mario Limonciello <mario.limonciello@dell.com>
    Reviewed-by: Pali Rohár <pali.rohar@gmail.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 21fd877300b01d25c5807c327848fdc7c813cf0e
Author: Sergio Paracuellos <sergio.paracuellos@gmail.com>
Date:   Sun Oct 6 20:10:32 2019 +0200

    staging: mt7621-pci: add quirks for 'E2' revision using 'soc_device_attribute'

    [ Upstream commit b483b4e4d3f6bfd5089b9e6dc9ba259879c6ce6f ]

    Depending on revision of the chip, reset lines are inverted. Make code
    more readable making use of 'soc_device_match' in driver probe function.

    Signed-off-by: Sergio Paracuellos <sergio.paracuellos@gmail.com>
    Link: https://lore.kernel.org/r/20191006181032.19112-1-sergio.paracuellos@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9ddede0a9fcda03e3c6ac429386f2f392640e72
Author: Andrii Nakryiko <andriin@fb.com>
Date:   Thu Oct 3 21:02:11 2019 -0700

    libbpf: Fix BTF-defined map's __type macro handling of arrays

    [ Upstream commit a53ba15d81995868651dd28a85d8045aef3d4e20 ]

    Due to a quirky C syntax of declaring pointers to array or function
    prototype, existing __type() macro doesn't work with map key/value types
    that are array or function prototype. One has to create a typedef first
    and use it to specify key/value type for a BPF map.  By using typeof(),
    pointer to type is now handled uniformly for all kinds of types. Convert
    one of self-tests as a demonstration.

    Signed-off-by: Andrii Nakryiko <andriin@fb.com>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Link: https://lore.kernel.org/bpf/20191004040211.2434033-1-andriin@fb.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d40471b727e8010fc5c9796dea36675e3d567c63
Author: Jiange Zhao <Jiange.Zhao@amd.com>
Date:   Thu Sep 19 13:22:59 2019 -0500
```

```
    drm/amdgpu/SRIOV: add navi12 pci id for SRIOV (v2)

    [ Upstream commit 57d4f3b7fd65b56f98b62817f27c461142c0bc2a ]

    Add Navi12 PCI id support.

    v2: flag as experimental for now (Alex)

    Signed-off-by: Jiange Zhao <Jiange.Zhao@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 791e3ec642cec12b090003d36ab1465571023c9b
Author: Sam McNally <sammc@chromium.org>
Date:   Tue Sep 17 15:49:33 2019 +1000

    ASoC: Intel: cht_bsw_rt5645: Add quirk for boards using pmc_plt_clk_0

    [ Upstream commit adebb11139029ddf1fba6f796c4a476f17eacddc ]

    As of commit 648e921888ad ("clk: x86: Stop marking clocks as
    CLK_IS_CRITICAL"), the cht_bsw_rt5645 driver needs to enable the clock
    it's using for the codec's mclk. It does this from commit 7735bce05a9c
    ("ASoC: Intel: boards: use devm_clk_get() unconditionally"), enabling
    pmc_plt_clk_3. However, Strago family Chromebooks use pmc_plt_clk_0 for
    the codec mclk, resulting in white noise with some digital microphones.
    Add a DMI-based quirk for Strago family Chromebooks to use pmc_plt_clk_0
    instead - mirroring the changes made to cht_bsw_max98090_ti in
    commit a182ecd3809c ("ASoC: intel: cht_bsw_max98090_ti: Add quirk for
    boards using pmc_plt_clk_0") and making use of the existing
    dmi_check_system() call and related infrastructure added in
    commit 22af29114eb4 ("ASoC: Intel: cht-bsw-rt5645: add quirks for
    SSP0/AIF1/AIF2 routing").

    Signed-off-by: Sam McNally <sammc@chromium.org>
    Acked-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20190917054933.209335-1-sammc@chromium.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48617aa63951879d9f144c041b67b8c0bf78265a
Author: Yauhen Kharuzhy <jekhor@gmail.com>
Date:   Tue Sep 17 00:15:36 2019 +0300

    extcon-intel-cht-wc: Don't reset USB data connection at probe

    [ Upstream commit e81b88932985c9134d410f4eaaaa9b81a3b4bd0c ]

    Intel Cherry Trail Whiskey Cove extcon driver connect USB data lines to
    PMIC at driver probing for further charger detection. This causes reset of
    USB data sessions and removing all devices from bus. If system was
    booted from Live CD or USB dongle, this makes system unusable.

    Check if USB ID pin is floating and re-route data lines in this case
    only, don't touch otherwise.

    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Yauhen Kharuzhy <jekhor@gmail.com>
    [cw00.choi: Clean-up the minor coding style]
    Signed-off-by: Chanwoo Choi <cw00.choi@samsung.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit de850e86d2e1a7702d9561f6fdaf317af82a3bd4
Author: Rodrigo Rivas Costa <rodrigorivascosta@gmail.com>
Date:   Tue Jan 7 20:48:13 2020 +0100

    HID: steam: Fix input device disappearing

    [ Upstream commit 20eee6e5af35d9586774e80b6e0b1850e7cc9899 ]

    The `connected` value for wired devices was not properly initialized,
    it must be set to `true` upon creation, because wired devices do not
    generate connection events.

    When a raw client (the Steam Client) uses the device, the input device
    is destroyed. Then, when the raw client finishes, it must be recreated.
    But since the `connected` variable was false this never happended.

    Signed-off-by: Rodrigo Rivas Costa <rodrigorivascosta@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b60d320dc3630a3042987fc229dcf1bd3a070512
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Tue Jan 7 21:43:59 2020 +0100

    atm: eni: fix uninitialized variable warning

    [ Upstream commit 30780d086a83332adcd9362281201cee7c3d9d19 ]

    With -O3, gcc has found an actual unintialized variable stored
    into an mmio register in two instances:

    drivers/atm/eni.c: In function 'discard':
    drivers/atm/eni.c:465:13: error: 'dma[1]' is used uninitialized in this function [-
Werror=uninitialized]
       writel(dma[i*2+1],eni_dev->rx_dma+dma_wr*8+4);
              ^
    drivers/atm/eni.c:465:13: error: 'dma[3]' is used uninitialized in this function [-
Werror=uninitialized]

    Change the code to always write zeroes instead.

    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f4691c14d449cd8d884ecbd2bb0e856216da2939
Author: Jiping Ma <jiping.ma2@windriver.com>
Date:   Tue Jan 7 14:34:00 2020 +0800

    stmmac: debugfs entry name is not be changed when udev rename device name.

    [ Upstream commit 481a7d154cbbd5ca355cc01cc8969876b240eded ]

    Add one notifier for udev changes net device name.
    Fixes: b6601323ef9e ("net: stmmac: debugfs entry name is not be changed when udev
rename")

    Signed-off-by: Jiping Ma <jiping.ma2@windriver.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d54646059b2755f8f42fc24c2ab0f6097216113
Author: Thomas Anderson <thomasanderson@google.com>
Date:   Mon Dec 2 13:47:13 2019 -0800
```

```
    drm/amd/display: Reduce HDMI pixel encoding if max clock is exceeded

    [ Upstream commit 840c90fce6c78bc6b2c4cb9e836d70985ed32066 ]

    For high-res (8K) or HFR (4K120) displays, using uncompressed pixel
    formats like YCbCr444 would exceed the bandwidth of HDMI 2.0, so the
    "interesting" modes would be disabled, leaving only low-res or low
    framerate modes.

    This change lowers the pixel encoding to 4:2:2 or 4:2:0 if the max TMDS
    clock is exceeded. Verified that 8K30 and 4K120 are now available and
    working with a Samsung Q900R over an HDMI 2.0b link from a Radeon 5700.

    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Thomas Anderson <thomasanderson@google.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a347d66cbec7503bda3c7ef551600bfa303b326e
Author: Qian Cai <cai@lca.pw>
Date:   Mon Jan 6 10:27:27 2020 -0500

    iommu/dma: fix variable 'cookie' set but not used

    [ Upstream commit 55817b340a31951d23d1692db45522560b1d20f9 ]

    The commit c18647900ec8 ("iommu/dma: Relax locking in
    iommu_dma_prepare_msi()") introduced a compliation warning,

    drivers/iommu/dma-iommu.c: In function 'iommu_dma_prepare_msi':
    drivers/iommu/dma-iommu.c:1206:27: warning: variable 'cookie' set but
    not used [-Wunused-but-set-variable]
      struct iommu_dma_cookie *cookie;
                               ^~~~~~

    Fixes: c18647900ec8 ("iommu/dma: Relax locking in iommu_dma_prepare_msi()")
    Signed-off-by: Qian Cai <cai@lca.pw>
    Acked-by: Robin Murphy <robin.murphy@arm.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fe14f9db5c948ab0def53fda4275d958c2d26fb
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Mon Jan 6 04:51:54 2020 +0300

    gpio: max77620: Add missing dependency on GPIOLIB_IRQCHIP

    [ Upstream commit c5706c7defc79de68a115b5536376298a8fef111 ]

    Driver fails to compile in a minimized kernel's configuration because of
    the missing dependency on GPIOLIB_IRQCHIP.

     error: ‘struct gpio_chip’ has no member named ‘irq’
       44 |   virq = irq_find_mapping(gpio->gpio_chip.irq.domain, offset);

    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Link: https://lore.kernel.org/r/20200106015154.12040-1-digetx@gmail.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dfe190e778625d38433dc19bf3dacbbe92fd1f08
Author: Krzysztof Kozlowski <krzk@kernel.org>
Date:   Sat Jan 4 15:31:43 2020 +0100

    net: wan: sdla: Fix cast from pointer to integer of different size
```

```
    [ Upstream commit 00c0688cecadbf7ac2f5b4cdb36d912a2d3f0cca ]

    Since net_device.mem_start is unsigned long, it should not be cast to
    int right before casting to pointer.  This fixes warning (compile
    testing on alpha architecture):

        drivers/net/wan/sdla.c: In function ‘sdla_transmit’:
        drivers/net/wan/sdla.c:711:13: warning:
            cast to pointer from integer of different size [-Wint-to-pointer-cast]

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

commit eea2b2bb502f4e025d7be664336ef6a46536810f
Author: Fenghua Yu <fenghua.yu@intel.com>
Date:   Thu Jan 2 13:27:06 2020 -0800

    drivers/net/b44: Change to non-atomic bit operations on pwol_mask

    [ Upstream commit f11421ba4af706cb4f5703de34fa77fba8472776 ]

    Atomic operations that span cache lines are super-expensive on x86
    (not just to the current processor, but also to other processes as all
    memory operations are blocked until the operation completes). Upcoming
    x86 processors have a switch to cause such operations to generate a #AC
    trap. It is expected that some real time systems will enable this mode
    in BIOS.

    In preparation for this, it is necessary to fix code that may execute
    atomic instructions with operands that cross cachelines because the #AC
    trap will crash the kernel.

    Since "pwol_mask" is local and never exposed to concurrency, there is
    no need to set bits in pwol_mask using atomic operations.

    Directly operate on the byte which contains the bit instead of using
    __set_bit() to avoid any big endian concern due to type cast to
    unsigned long in __set_bit().

    Suggested-by: Peter Zijlstra <peterz@infradead.org>
    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Signed-off-by: Tony Luck <tony.luck@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2ae3ce351be6a12bba55870ac83b8d068fe5d30
Author: Liran Alon <liran.alon@oracle.com>
Date:   Fri Jan 3 18:44:59 2020 +0200

    net: Google gve: Remove dma_wmb() before ringing doorbell

    [ Upstream commit b54ef37b1ce892fdf6b632d566246d2f2f539910 ]

    Current code use dma_wmb() to ensure Rx/Tx descriptors are visible
    to device before writing to doorbell.

    However, these dma_wmb() are wrong and unnecessary. Therefore,
    they should be removed.

    iowrite32be() called from gve_rx_write_doorbell()/gve_tx_put_doorbell()
    should guaratee that all previous writes to WB/UC memory is visible to
    device before the write done by iowrite32be().
```

```
    E.g. On ARM64, iowrite32be() calls __iowmb() which expands to dma_wmb()
    and only then calls __raw_writel().

    Reviewed-by: Si-Wei Liu <si-wei.liu@oracle.com>
    Signed-off-by: Liran Alon <liran.alon@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7db4e6c728cbb4caf6708b0181bc11763d1e89a7
Author: wuxu.wu <wuxu.wu@huawei.com>
Date:   Wed Jan 1 11:39:41 2020 +0800

    spi: spi-dw: Add lock protect dw_spi rx/tx to prevent concurrent calls

    [ Upstream commit 19b61392c5a852b4e8a0bf35aecb969983c5932d ]

    dw_spi_irq() and dw_spi_transfer_one concurrent calls.

    I find a panic in dw_writer(): txw = *(u8 *)(dws->tx), when dw->tx==null,
    dw->len==4, and dw->tx_end==1.

    When tpm driver's message overtime dw_spi_irq() and dw_spi_transfer_one
    may concurrent visit dw_spi, so I think dw_spi structure lack of protection.

    Otherwise dw_spi_transfer_one set dw rx/tx buffer and then open irq,
    store dw rx/tx instructions and other cores handle irq load dw rx/tx
    instructions may out of order.

            [ 1025.321302] Call trace:
            ...
            [ 1025.321319]  __crash_kexec+0x98/0x148
            [ 1025.321323]  panic+0x17c/0x314
            [ 1025.321329]  die+0x29c/0x2e8
            [ 1025.321334]  die_kernel_fault+0x68/0x78
            [ 1025.321337]  __do_kernel_fault+0x90/0xb0
            [ 1025.321346]  do_page_fault+0x88/0x500
            [ 1025.321347]  do_translation_fault+0xa8/0xb8
            [ 1025.321349]  do_mem_abort+0x68/0x118
            [ 1025.321351]  el1_da+0x20/0x8c
            [ 1025.321362]  dw_writer+0xc8/0xd0
            [ 1025.321364]  interrupt_transfer+0x60/0x110
            [ 1025.321365]  dw_spi_irq+0x48/0x70
            ...

    Signed-off-by: wuxu.wu <wuxu.wu@huawei.com>
    Link: https://lore.kernel.org/r/1577849981-31489-1-git-send-email-wuxu.wu@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db10c64764eb897258c386fbab8cd2aec1f6c8c9
Author: Sean Nyekjaer <sean@geanix.com>
Date:   Wed Dec 11 14:58:51 2019 +0100

    can: tcan4x5x: tcan4x5x_parse_config(): reset device before register access

    [ Upstream commit c3083124e6a1c0d6cd4fe3b3f627b488bd3b10c4 ]

    It's a good idea to reset a ip-block/spi device before using it, this
    patch will reset the device.

    And a generic reset function if needed elsewhere.

    Signed-off-by: Sean Nyekjaer <sean@geanix.com>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 68d199fc8e668263b5586c4bb4446e33d5f347c6
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Mon Dec 16 10:24:32 2019 -0600

    usb: musb: jz4740: Silence error if code is -EPROBE_DEFER

    [ Upstream commit ce03cbcb4b4fd2a3817f32366001f1ca45d213b8 ]

    Avoid printing any error message if the error code is -EPROBE_DEFER.

    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Bin Liu <b-liu@ti.com>
    Link: https://lore.kernel.org/r/20191216162432.1256-1-b-liu@ti.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ecda03a5c13086b0dd1d0cf0c5fe58d9c2afce9
Author: Russell King <rmk+kernel@armlinux.org.uk>
Date:   Fri Nov 29 14:51:35 2019 +0000

    watchdog: orion: fix platform_get_irq() complaints

    [ Upstream commit dcbce5fbcc69bf2553f650004aad44bf390eca73 ]

    Fix:

    orion_wdt f1020300.watchdog: IRQ index 1 not found

    which is caused by platform_get_irq() now complaining when optional
    IRQs are not found.  Neither interrupt for orion is required, so
    make them both optional.

    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/E1iahcN-0000AT-Co@rmk-PC.armlinux.org.uk
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b39f38f0ebed1d1da859aff71dd483383922ca8a
Author: Andreas Kemnade <andreas@kemnade.info>
Date:   Fri Dec 13 22:48:02 2019 +0100

    watchdog: rn5t618_wdt: fix module aliases

    [ Upstream commit a76dfb859cd42df6e3d1910659128ffcd2fb6ba2 ]

    Platform device aliases were missing so module autoloading
    did not work.

    Signed-off-by: Andreas Kemnade <andreas@kemnade.info>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20191213214802.22268-1-andreas@kemnade.info
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 27757bfb8c72d88371bde6f1372acfc35d52fc55
Author: David Engraf <david.engraf@sysgo.com>
Date:   Wed Nov 27 09:46:17 2019 +0100

    watchdog: max77620_wdt: fix potential build errors

    [ Upstream commit da9e3f4e30a53cd420cf1e6961c3b4110f0f21f0 ]
```

```
    max77620_wdt uses watchdog core functions. Enable CONFIG_WATCHDOG_CORE
    to fix potential build errors.

    Signed-off-by: David Engraf <david.engraf@sysgo.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20191127084617.16937-1-david.engraf@sysgo.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5399714dc2a1103e65e1fd37b8b6f6a51eb875a2
Author: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
Date:   Fri Dec 27 13:40:57 2019 -0800

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

    [ Upstream commit 1479a82d82df68dfac29c72c774cb8bdc17d4eb1 ]

    Added Tiger Lake PCI device ID to the supported device list.

    Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e5900bc2225395677d299e5385c5fa9ea2840112
Author: Tony Lindgren <tony@atomide.com>
Date:   Sun Dec 22 10:17:02 2019 -0800

    phy: cpcap-usb: Prevent USB line glitches from waking up modem

    [ Upstream commit 63078b6ba09e842f09df052c5728857389fddcd2 ]

    The micro-USB connector on Motorola Mapphone devices can be muxed between
    the SoC and the mdm6600 modem. But even when used for the SoC, configuring
    the PHY with ID pin grounded will wake up the modem from idle state. Looks
    like the issue is probably caused by line glitches.

    We can prevent the glitches by using a previously unknown mode of the
    GPIO mux to prevent the USB lines from being connected to the moden while
    configuring the USB PHY, and enable the USB lines after configuring the
    PHY.

    Note that this only prevents waking up mdm6600 as regular USB A-host mode,
    and does not help when connected to a lapdock. The lapdock specific issue
    still needs to be debugged separately.

    Cc: Merlijn Wajer <merlijn@wizzup.org>
    Cc: Pavel Machek <pavel@ucw.cz>
    Cc: Sebastian Reichel <sre@kernel.org>
    Acked-by: Pavel Machek <pavel@ucw.cz>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 98aedb8935551b1e8934e3482c6b8e6884da8fba
Author: Dragos Tarcatu <dragos_tarcatu@mentor.com>
Date:   Wed Dec 4 15:04:47 2019 -0600

    ASoC: topology: Prevent use-after-free in snd_soc_get_pcm_runtime()

    [ Upstream commit dd836ddf4e4e1c7f1eb2ae44783ccd70872ef24e ]

    remove_link() is currently calling snd_soc_remove_dai_link() after
    it has already freed the memory for the link name. But this is later
```

```
    read from snd_soc_get_pcm_runtime() causing a KASAN use-after-free
    warning. Reorder the cleanups to fix this issue.

    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Dragos Tarcatu <dragos_tarcatu@mentor.com>
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
    Link: https://lore.kernel.org/r/20191204210447.11701-4-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68b74847c8bcec2728332c8929ae3d3cf0e83338
Author: Chuhong Yuan <hslester96@gmail.com>
Date:   Tue Dec 3 19:13:03 2019 +0800

    ASoC: fsl_audmix: add missed pm_runtime_disable

    [ Upstream commit 77fffa742285f2b587648d6c72b5c705633f146f ]

    The driver forgets to call pm_runtime_disable in probe failure
    and remove.
    Add the missed calls to fix it.

    Signed-off-by: Chuhong Yuan <hslester96@gmail.com>
    Acked-by: Nicolin Chen <nicoleotsuka@gmail.com>
    Link: https://lore.kernel.org/r/20191203111303.12933-1-hslester96@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a0636e211f8ed74dddcc4e966de1b9d70aaf4ccc
Author: Bjorn Andersson <bjorn.andersson@linaro.org>
Date:   Fri Dec 20 15:47:15 2019 +0530

    phy: qcom-qmp: Increase PHY ready timeout

    [ Upstream commit cd217ee6867d285ceecd610fa1006975d5c683fa ]

    It's typical for the QHP PHY to take slightly above 1ms to initialize,
    so increase the timeout of the PHY ready check to 10ms - as already done
    in the downstream PCIe driver.

    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Tested-by: Evan Green <evgreen@chromium.org>
    Tested-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b826ffe45a27b28d78a540afdb795c8f2c41b08b
Author: Pan Zhang <zhangpan26@huawei.com>
Date:   Thu Dec 19 10:45:30 2019 +0800

    drivers/hid/hid-multitouch.c: fix a possible null pointer access.

    [ Upstream commit 306d5acbfc66e7cccb4d8f91fc857206b8df80d1 ]

    1002     if ((quirks & MT_QUIRK_IGNORE_DUPLICATES) && mt) {
    1003         struct input_mt_slot *i_slot = &mt->slots[slotnum];
    1004
    1005         if (input_mt_is_active(i_slot) &&
    1006             input_mt_is_used(mt, i_slot))
    1007             return -EAGAIN;
    1008     }
```

```
    We previously assumed 'mt' could be null (see line 1002).

    The following situation is similar, so add a judgement.

    Signed-off-by: Pan Zhang <zhangpan26@huawei.com>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32a369776bd54a8b3fee52883eb40a9368cac2ae
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Tue Dec 17 18:05:17 2019 -0600

    ASoC: SOF: Intel: hda: hda-dai: fix oops on hda_link .hw_free

    [ Upstream commit 921162c81a089aa2f442103290f1af9ba281fc9f ]

    When the PCM_PARAM IPC fails while configuring the FE, the kernel
    oopses in the HDaudio link DMA .hw_free operation. The root cause is a
    NULL dma_data since the BE .hw_params was never called by the SOC
    core.

    This error can also happen if the HDaudio link DMA configuration IPC
    fails in the BE .hw_params.

    This patches makes sure the dma_data is properly saved in .hw_params,
    and tested before being use in hw_free.

    GitHub issue: https://github.com/thesofproject/linux/issues/1417

    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Reviewed-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20191218000518.5830-3-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eedf205f026757f0f9cd2128e4ff73ea34b17d55
Author: Kai Vehmanen <kai.vehmanen@linux.intel.com>
Date:   Tue Dec 17 18:05:16 2019 -0600

    ASoC: SOF: fix fault at driver unload after failed probe

    [ Upstream commit b06e46427f987bf83dcb6a69516b57276eb8ec0c ]

    If sof_machine_check() fails during driver probe, the IPC
    state is not initialized and this will lead to a NULL
    dereference at driver unload. Example log is as follows:

    [ 1535.980630] sof-audio-pci 0000:00:1f.3: error: no matching ASoC machine driver found
- aborting probe
    [ 1535.980631] sof-audio-pci 0000:00:1f.3: error: failed to get machine info -19
    [ 1535.980632] sof-audio-pci 0000:00:1f.3: error: sof_probe_work failed err: -19
    [ 1550.798373] BUG: kernel NULL pointer dereference, address: 0000000000000008
    ...
    [ 1550.798393] Call Trace:
    [ 1550.798397]  snd_sof_ipc_free+0x15/0x30 [snd_sof]
    [ 1550.798399]  snd_sof_device_remove+0x29/0xa0 [snd_sof]
    [ 1550.798400]  sof_pci_remove+0x10/0x30 [snd_sof_pci]

    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20191218000518.5830-2-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
```

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

commit 8cc2de41a5293efc3ae050695bcc7edf5132997b
Author: Jason Gerecke <killertofu@gmail.com>
Date:   Mon Dec 16 10:18:57 2019 -0800

    HID: wacom: Recognize new MobileStudio Pro PID

    [ Upstream commit fe4e940f0f91b4a506f048b42e00386f5ad322b6 ]

    A new PID is in use for repaired MobileStudio Pro devices. Add it to the
    list of devices that need special-casing in wacom_wac_pad_event.

    Signed-off-by: Jason Gerecke <jason.gerecke@wacom.com>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b55f9e9760c85e5ccbef2f693caeba8a6edb107b
Author: Even Xu <even.xu@intel.com>
Date:   Tue Dec 17 09:11:18 2019 +0800

    HID: intel-ish-hid: ipc: add CMP device id

    [ Upstream commit abb33ee8a8c0d146b4b2c52937dc86a15ec92d05 ]

    Add Comet Lake H into ishtp support list.

    Signed-off-by: Even Xu <even.xu@intel.com>
    Acked-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22aa1db061d216d5a63a1a5fb12f698280edfdf3
Author: Pavel Balan <admin@kryma.net>
Date:   Wed Nov 27 03:23:29 2019 +0000

    HID: Add quirk for incorrect input length on Lenovo Y720

    [ Upstream commit fd0913768701612fc2b8ab9c8a5c019133e8d978 ]

    Apply it to the Lenovo Y720 gaming laptop I2C peripheral then.

    This fixes dmesg being flooded with errors visible on un-suspend
    in Linux Mint 19 Cinnamon.

    Example of error log:

    <...>
    [    4.326588] i2c_hid i2c-ITE33D1:00: i2c_hid_get_input: incomplete report (2/4)
    [    4.326845] i2c_hid i2c-ITE33D1:00: i2c_hid_get_input: incomplete report (2/4)
    [    4.327095] i2c_hid i2c-ITE33D1:00: i2c_hid_get_input: incomplete report (2/4)
    [    4.327341] i2c_hid i2c-ITE33D1:00: i2c_hid_get_input: incomplete report (2/4)
    [    4.327609] i2c_hid i2c-ITE33D1:00: i2c_hid_get_input: incomplete report (2/4)
    <...>

    Example of fixed log (debug on)

    <...>
    [ 3731.333183] i2c_hid i2c-ITE33D1:00: input: 02 00
    [ 3731.333581] i2c_hid i2c-ITE33D1:00: input: 02 00
    [ 3731.333842] i2c_hid i2c-ITE33D1:00: input: 02 00
    [ 3731.334107] i2c_hid i2c-ITE33D1:00: input: 02 00
    [ 3731.334367] i2c_hid i2c-ITE33D1:00: input: 02 00
    <...>
```

```
    [jkosina@suse.cz: rebase onto more recent codebase]
    Signed-off-by: Pavel Balan <admin@kryma.net>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 854f39866e2f1b64eea74b01ebca276ff68b9269
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Oct 24 13:59:09 2019 +0200

    HID: asus: Ignore Asus vendor-page usage-code 0xff events

    [ Upstream commit c07a0254c89e4bb69ca781cd488baa5b628e2754 ]

    At least on a T100HA an Asus vendor-page usage-code 0xff event is send on
    every suspend and again on resume, resulting in the following warning:

    asus 0003:0B05:1807.0002: Unmapped Asus vendor usagepage code 0xff

    being logged twice on every suspend/resume.

    This commit silences the "Unmapped Asus vendor usagepage code ..."
    warning for usage-code 0xff to avoid these warnings being logged.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 49de93b4018b20bb5020d3a2b8d483ca28fc8a05
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Tue Nov 19 15:57:11 2019 +0100

    HID: ite: Add USB id match for Acer SW5-012 keyboard dock

    [ Upstream commit 8f18eca9ebc57d6b150237033f6439242907e0ba ]

    The Acer SW5-012 2-in-1 keyboard dock uses a Synaptics S91028 touchpad
    which is connected to an ITE 8595 USB keyboard controller chip.

    This keyboard has the same quirk for its rfkill / airplane mode hotkey as
    other keyboards with the ITE 8595 chip, it only sends a single release
    event when pressed and released, it never sends a press event.

    This commit adds this keyboards USB id to the hid-ite id-table, fixing
    the rfkill key not working on this keyboard.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05c33c23d0b651c811487855b38a2693b39547b7
Author: Priit Laes <plaes@plaes.org>
Date:   Sun Dec 1 00:22:09 2019 +0200

    HID: Add quirk for Xin-Mo Dual Controller

    [ Upstream commit c62f7cd8ed066a93a243643ebf57ca99f754388e ]

    Without the quirk, joystick shows up as single controller
    for both first and second player pads/pins.

    Signed-off-by: Priit Laes <plaes@plaes.org>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae9484c6fb6024e70eee942cecfea5022ea2b741
```

```
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Tue Nov 26 17:54:09 2019 -0800

    arc: eznps: fix allmodconfig kconfig warning

    [ Upstream commit 1928b36cfa4df1aeedf5f2644d0c33f3a1fcfd7b ]

    Fix kconfig warning for arch/arc/plat-eznps/Kconfig allmodconfig:

    WARNING: unmet direct dependencies detected for CLKSRC_NPS
      Depends on [n]: GENERIC_CLOCKEVENTS [=y] && !PHYS_ADDR_T_64BIT [=y]
      Selected by [y]:
      - ARC_PLAT_EZNPS [=y]

    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Vineet Gupta <vgupta@synopsys.com>
    Cc: Ofer Levi <oferle@mellanox.com>
    Cc: linux-snps-arc@lists.infradead.org
    Signed-off-by: Vineet Gupta <vgupta@synopsys.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f5adacf21c6e016f5c14701e0a1605f6a941b285
Author: Aaron Ma <aaron.ma@canonical.com>
Date:   Wed Dec 4 20:42:07 2019 +0800

    HID: multitouch: Add LG MELF0410 I2C touchscreen support

    [ Upstream commit 348b80b273fbf4ce2a307f9e38eadecf37828cad ]

    Add multitouch support for LG MELF I2C touchscreen.
    Apply the same workaround as LG USB touchscreen.

    Signed-off-by: Aaron Ma <aaron.ma@canonical.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b576b27f7ba4597aab4f222be108258aee9e5c07
Author: David Howells <dhowells@redhat.com>
Date:   Fri Jan 24 23:08:04 2020 +0000

    rxrpc: Fix use-after-free in rxrpc_receive_data()

    [ Upstream commit 122d74fac84204b9a98263636f6f9a3b2e665639 ]

    The subpacket scanning loop in rxrpc_receive_data() references the
    subpacket count in the private data part of the sk_buff in the loop
    termination condition.  However, when the final subpacket is pasted into
    the ring buffer, the function is no longer has a ref on the sk_buff and
    should not be looking at sp->* any more.  This point is actually marked in
    the code when skb is cleared (but sp is not - which is an error).

    Fix this by caching sp->nr_subpackets in a local variable and using that
    instead.

    Also clear 'sp' to catch accesses after that point.

    This can show up as an oops in rxrpc_get_skb() if sp->nr_subpackets gets
    trashed by the sk_buff getting freed and reused in the meantime.

    Fixes: e2de6c404898 ("rxrpc: Use info in skbuff instead of reparsing a jumbo packet")
    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 d3da5e90665a103bcd067a08be174bb58e940c3e
```

```
Author: Stephen Worley <sworley@cumulusnetworks.com>
Date:   Fri Jan 24 16:53:27 2020 -0500

    net: include struct nhmsg size in nh nlmsg size

    [ Upstream commit f9e95555757915fc194288862d2978e370fe316b ]

    Include the size of struct nhmsg size when calculating
    how much of a payload to allocate in a new netlink nexthop
    notification message.

    Without this, we will fail to fill the skbuff at certain nexthop
    group sizes.

    You can reproduce the failure with the following iproute2 commands:

    ip link add dummy1 type dummy
    ip link add dummy2 type dummy
    ip link add dummy3 type dummy
    ip link add dummy4 type dummy
    ip link add dummy5 type dummy
    ip link add dummy6 type dummy
    ip link add dummy7 type dummy
    ip link add dummy8 type dummy
    ip link add dummy9 type dummy
    ip link add dummy10 type dummy
    ip link add dummy11 type dummy
    ip link add dummy12 type dummy
    ip link add dummy13 type dummy
    ip link add dummy14 type dummy
    ip link add dummy15 type dummy
    ip link add dummy16 type dummy
    ip link add dummy17 type dummy
    ip link add dummy18 type dummy
    ip link add dummy19 type dummy

    ip ro add 1.1.1.1/32 dev dummy1
    ip ro add 1.1.1.2/32 dev dummy2
    ip ro add 1.1.1.3/32 dev dummy3
    ip ro add 1.1.1.4/32 dev dummy4
    ip ro add 1.1.1.5/32 dev dummy5
    ip ro add 1.1.1.6/32 dev dummy6
    ip ro add 1.1.1.7/32 dev dummy7
    ip ro add 1.1.1.8/32 dev dummy8
    ip ro add 1.1.1.9/32 dev dummy9
    ip ro add 1.1.1.10/32 dev dummy10
    ip ro add 1.1.1.11/32 dev dummy11
    ip ro add 1.1.1.12/32 dev dummy12
    ip ro add 1.1.1.13/32 dev dummy13
    ip ro add 1.1.1.14/32 dev dummy14
    ip ro add 1.1.1.15/32 dev dummy15
    ip ro add 1.1.1.16/32 dev dummy16
    ip ro add 1.1.1.17/32 dev dummy17
    ip ro add 1.1.1.18/32 dev dummy18
    ip ro add 1.1.1.19/32 dev dummy19

    ip next add id 1 via 1.1.1.1 dev dummy1
    ip next add id 2 via 1.1.1.2 dev dummy2
    ip next add id 3 via 1.1.1.3 dev dummy3
    ip next add id 4 via 1.1.1.4 dev dummy4
    ip next add id 5 via 1.1.1.5 dev dummy5
    ip next add id 6 via 1.1.1.6 dev dummy6
    ip next add id 7 via 1.1.1.7 dev dummy7
    ip next add id 8 via 1.1.1.8 dev dummy8
    ip next add id 9 via 1.1.1.9 dev dummy9
```

```
    ip next add id 10 via 1.1.1.10 dev dummy10
    ip next add id 11 via 1.1.1.11 dev dummy11
    ip next add id 12 via 1.1.1.12 dev dummy12
    ip next add id 13 via 1.1.1.13 dev dummy13
    ip next add id 14 via 1.1.1.14 dev dummy14
    ip next add id 15 via 1.1.1.15 dev dummy15
    ip next add id 16 via 1.1.1.16 dev dummy16
    ip next add id 17 via 1.1.1.17 dev dummy17
    ip next add id 18 via 1.1.1.18 dev dummy18
    ip next add id 19 via 1.1.1.19 dev dummy19

    ip next add id 1111 group 1/2/3/4/5/6/7/8/9/10/11/12/13/14/15/16/17/18/19
    ip next del id 1111

    Fixes: 430a049190de ("nexthop: Add support for nexthop groups")
    Signed-off-by: Stephen Worley <sworley@cumulusnetworks.com>
    Reviewed-by: David Ahern <dsahern@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 73226c74746d4906ba1d591f4bca8f6253de14d7
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sat Jan 25 22:18:47 2020 +0100

    mlxsw: minimal: Fix an error handling path in 'mlxsw_m_port_create()'

    [ Upstream commit 6dd4b4f3936e17fedea1308bc70e9716f68bf232 ]

    An 'alloc_etherdev()' called is not ballanced by a corresponding
    'free_netdev()' call in one error handling path.

    Slighly reorder the error handling code to catch the missed case.

    Fixes: c100e47caa8e ("mlxsw: minimal: Add ethtool support")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Reviewed-by: Ido Schimmel <idosch@mellanox.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3c8c966cc7182b3975ef955a1258ec9e30dc3057
Author: Willem de Bruijn <willemb@google.com>
Date:   Mon Jan 27 15:40:31 2020 -0500

    udp: segment looped gso packets correctly

    [ Upstream commit 6cd021a58c18a1731f7e47f83e172c0c302d65e5 ]

    Multicast and broadcast packets can be looped from egress to ingress
    pre segmentation with dev_loopback_xmit. That function unconditionally
    sets ip_summed to CHECKSUM_UNNECESSARY.

    udp_rcv_segment segments gso packets in the udp rx path. Segmentation
    usually executes on egress, and does not expect packets of this type.
    __udp_gso_segment interprets !CHECKSUM_PARTIAL as CHECKSUM_NONE. But
    the offsets are not correct for gso_make_checksum.

    UDP GSO packets are of type CHECKSUM_PARTIAL, with their uh->check set
    to the correct pseudo header checksum. Reset ip_summed to this type.
    (CHECKSUM_PARTIAL is allowed on ingress, see comments in skbuff.h)

    Reported-by: syzbot <syzkaller@googlegroups.com>
    Fixes: cf329aa42b66 ("udp: cope with UDP GRO packet misdirection")
    Signed-off-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 87c681943725dad8a1fbd7f212c49779f213cbf1
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sat Jan 25 12:48:51 2020 +0100

    net: socionext: fix xdp_result initialization in netsec_process_rx

    [ Upstream commit 02758cb6dac31a2b4bd9e535cffbe718acd46404 ]

    Fix xdp_result initialization in netsec_process_rx in order to not
    increase rx counters if there is no bpf program attached to the xdp hook
    and napi_gro_receive returns GRO_DROP

    Fixes: ba2b232108d3c ("net: netsec: add XDP support")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Acked-by: Jesper Dangaard Brouer <brouer@redhat.com>
    Acked-by: Ilias Apalodimas <ilias.apalodimas@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b7eda225ebc1c2faa793ac2343c6d34b3b645a66
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sat Jan 25 12:48:50 2020 +0100

    net: socionext: fix possible user-after-free in netsec_process_rx

    [ Upstream commit b5e82e3c89c78ee0407ea8e8087af5519b6c7bae ]

    Fix possible use-after-free in in netsec_process_rx that can occurs if
    the first packet is sent to the normal networking stack and the
    following one is dropped by the bpf program attached to the xdp hook.
    Fix the issue defining the skb pointer in the 'budget' loop

    Fixes: ba2b232108d3c ("net: netsec: add XDP support")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Acked-by: Jesper Dangaard Brouer <brouer@redhat.com>
    Acked-by: Ilias Apalodimas <ilias.apalodimas@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 717abad8646011b6eba7a19506f354905a461adf
Author: Cong Wang <xiyou.wangcong@gmail.com>
Date:   Thu Jan 23 17:27:08 2020 -0800

    net_sched: walk through all child classes in tc_bind_tclass()

    [ Upstream commit 760d228e322e99cdf6d81b4b60a268b8f13cf67a ]

    In a complex TC class hierarchy like this:

    tc qdisc add dev eth0 root handle 1:0 cbq bandwidth 100Mbit         \
      avpkt 1000 cell 8
    tc class add dev eth0 parent 1:0 classid 1:1 cbq bandwidth 100Mbit  \
      rate 6Mbit weight 0.6Mbit prio 8 allot 1514 cell 8 maxburst 20      \
      avpkt 1000 bounded

    tc filter add dev eth0 parent 1:0 protocol ip prio 1 u32 match ip \
      sport 80 0xffff flowid 1:3
    tc filter add dev eth0 parent 1:0 protocol ip prio 1 u32 match ip \
      sport 25 0xffff flowid 1:4

    tc class add dev eth0 parent 1:1 classid 1:3 cbq bandwidth 100Mbit  \
      rate 5Mbit weight 0.5Mbit prio 5 allot 1514 cell 8 maxburst 20      \
      avpkt 1000
    tc class add dev eth0 parent 1:1 classid 1:4 cbq bandwidth 100Mbit  \
```

```
       rate 3Mbit weight 0.3Mbit prio 5 allot 1514 cell 8 maxburst 20      \
       avpkt 1000

     where filters are installed on qdisc 1:0, so we can't merely
     search from class 1:1 when creating class 1:3 and class 1:4. We have
     to walk through all the child classes of the direct parent qdisc.
     Otherwise we would miss filters those need reverse binding.

     Fixes: 07d79fc7d94e ("net_sched: add reverse binding for tc class")
     Cc: Jamal Hadi Salim <jhs@mojatatu.com>
     Cc: Jiri Pirko <jiri@resnulli.us>
     Signed-off-by: Cong Wang <xiyou.wangcong@gmail.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 55ec468d309373845e51ab4bcc8f490321b55f70
  Author: Cong Wang <xiyou.wangcong@gmail.com>
  Date:   Thu Jan 23 16:26:18 2020 -0800

     net_sched: fix ops->bind_class() implementations

     [ Upstream commit 2e24cd755552350b94a7617617c6877b8cbcb701 ]

     The current implementations of ops->bind_class() are merely
     searching for classid and updating class in the struct tcf_result,
     without invoking either of cl_ops->bind_tcf() or
     cl_ops->unbind_tcf(). This breaks the design of them as qdisc's
     like cbq use them to count filters too. This is why syzbot triggered
     the warning in cbq_destroy_class().

     In order to fix this, we have to call cl_ops->bind_tcf() and
     cl_ops->unbind_tcf() like the filter binding path. This patch does
     so by refactoring out two helper functions __tcf_bind_filter()
     and __tcf_unbind_filter(), which are lockless and accept a Qdisc
     pointer, then teaching each implementation to call them correctly.

     Note, we merely pass the Qdisc pointer as an opaque pointer to
     each filter, they only need to pass it down to the helper
     functions without understanding it at all.

     Fixes: 07d79fc7d94e ("net_sched: add reverse binding for tc class")
     Reported-and-tested-by: syzbot+0a0596220218fcb603a8@syzkaller.appspotmail.com
     Reported-and-tested-by: syzbot+63bdb6006961d8c917c6@syzkaller.appspotmail.com
     Cc: Jamal Hadi Salim <jhs@mojatatu.com>
     Cc: Jiri Pirko <jiri@resnulli.us>
     Signed-off-by: Cong Wang <xiyou.wangcong@gmail.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 2d0b854393afaea3a18737ace523be38fdc7e438
  Author: Eric Dumazet <edumazet@google.com>
  Date:   Fri Jan 24 14:57:20 2020 -0800

     net_sched: ematch: reject invalid TCF_EM_SIMPLE

     [ Upstream commit 55cd9f67f1e45de8517cdaab985fb8e56c0bc1d8 ]

     It is possible for malicious userspace to set TCF_EM_SIMPLE bit
     even for matches that should not have this bit set.

     This can fool two places using tcf_em_is_simple()

     1) tcf_em_tree_destroy() -> memory leak of em->data
        if ops->destroy() is NULL
```

```
    2) tcf_em_tree_dump() wrongly report/leak 4 low-order bytes
       of a kernel pointer.

    BUG: memory leak
    unreferenced object 0xffff888121850a40 (size 32):
      comm "syz-executor927", pid 7193, jiffies 4294941655 (age 19.840s)
      hex dump (first 32 bytes):
        00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 00  ................
        00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
      backtrace:
        [<00000000f67036ea>] kmemleak_alloc_recursive include/linux/kmemleak.h:43 [inline]
        [<00000000f67036ea>] slab_post_alloc_hook mm/slab.h:586 [inline]
        [<00000000f67036ea>] slab_alloc mm/slab.c:3320 [inline]
        [<00000000f67036ea>] __do_kmalloc mm/slab.c:3654 [inline]
        [<00000000f67036ea>] __kmalloc_track_caller+0x165/0x300 mm/slab.c:3671
        [<00000000fab0cc8e>] kmemdup+0x27/0x60 mm/util.c:127
        [<00000000d9992e0a>] kmemdup include/linux/string.h:453 [inline]
        [<00000000d9992e0a>] em_nbyte_change+0x5b/0x90 net/sched/em_nbyte.c:32
        [<000000007e04f711>] tcf_em_validate net/sched/ematch.c:241 [inline]
        [<000000007e04f711>] tcf_em_tree_validate net/sched/ematch.c:359 [inline]
        [<000000007e04f711>] tcf_em_tree_validate+0x332/0x46f net/sched/ematch.c:300
        [<000000007a769204>] basic_set_parms net/sched/cls_basic.c:157 [inline]
        [<000000007a769204>] basic_change+0x1d7/0x5f0 net/sched/cls_basic.c:219
        [<00000000e57a5997>] tc_new_tfilter+0x566/0xf70 net/sched/cls_api.c:2104
        [<0000000074b68559>] rtnetlink_rcv_msg+0x3b2/0x4b0 net/core/rtnetlink.c:5415
        [<00000000b7fe53fb>] netlink_rcv_skb+0x61/0x170 net/netlink/af_netlink.c:2477
        [<00000000e83a40d0>] rtnetlink_rcv+0x1d/0x30 net/core/rtnetlink.c:5442
        [<00000000d62ba933>] netlink_unicast_kernel net/netlink/af_netlink.c:1302 [inline]
        [<00000000d62ba933>] netlink_unicast+0x223/0x310 net/netlink/af_netlink.c:1328
        [<0000000088070f72>] netlink_sendmsg+0x2c0/0x570 net/netlink/af_netlink.c:1917
        [<00000000f70b15ea>] sock_sendmsg_nosec net/socket.c:639 [inline]
        [<00000000f70b15ea>] sock_sendmsg+0x54/0x70 net/socket.c:659
        [<00000000ef95a9be>] ____sys_sendmsg+0x2d0/0x300 net/socket.c:2330
        [<00000000b650f1ab>] ___sys_sendmsg+0x8a/0xd0 net/socket.c:2384
        [<0000000055bfa74a>] __sys_sendmsg+0x80/0xf0 net/socket.c:2417
        [<000000002abac183>] __do_sys_sendmsg net/socket.c:2426 [inline]
        [<000000002abac183>] __se_sys_sendmsg net/socket.c:2424 [inline]
        [<000000002abac183>] __x64_sys_sendmsg+0x23/0x30 net/socket.c:2424

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: syzbot+03c4738ed29d5d366ddf@syzkaller.appspotmail.com
    Cc: Cong Wang <xiyou.wangcong@gmail.com>
    Acked-by: Cong Wang <xiyou.wangcong@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 79c13c063720b3b5e91ee4dd9be506149ce6a5d9
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:26 2019 +0100

    zd1211rw: fix storage endpoint lookup

    commit 2d68bb2687abb747558b933e80845ff31570a49c upstream.

    Make sure to use the current alternate setting when verifying the
    storage interface descriptors to avoid submitting an URB to an invalid
    endpoint.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: a1030e92c150 ("[PATCH] zd1211rw: Convert installer CDROM device into WLAN
device")
    Cc: stable <stable@vger.kernel.org>     # 2.6.19
```

```
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 226ed25115b7fc23ef57292ab690976a4b605e09
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:24 2019 +0100

    rtl8xxxu: fix interface sanity check

    commit 39a4281c312f2d226c710bc656ce380c621a2b16 upstream.

    Make sure to use the current alternate setting when verifying the
    interface descriptors to avoid binding to an invalid interface.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: 26f1fad29ad9 ("New driver: rtl8xxxu (mac80211)")
    Cc: stable <stable@vger.kernel.org>     # 4.4
    Cc: Jes Sorensen <Jes.Sorensen@redhat.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 670a39aeb07bee6a1e61c2f3857f8b3cab5c49b8
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:22 2019 +0100

    brcmfmac: fix interface sanity check

    commit 3428fbcd6e6c0850b1a8b2a12082b7b2aabb3da3 upstream.

    Make sure to use the current alternate setting when verifying the
    interface descriptors to avoid binding to an invalid interface.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: 71bb244ba2fd ("brcm80211: fmac: add USB support for bcm43235/6/8 chipsets")
    Cc: stable <stable@vger.kernel.org>     # 3.4
    Cc: Arend van Spriel <arend@broadcom.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b896fe1da14f1e32a0624b628f9309b1f7b606ff
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:20 2019 +0100

    ath9k: fix storage endpoint lookup

    commit 0ef332951e856efa89507cdd13ba8f4fb8d4db12 upstream.

    Make sure to use the current alternate setting when verifying the
    storage interface descriptors to avoid submitting an URB to an invalid
    endpoint.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: 36bcce430657 ("ath9k_htc: Handle storage devices")
    Cc: stable <stable@vger.kernel.org>     # 2.6.39
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
```

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

commit e088841970826a841e0a48ec3b3d57c807c2cb2f
Author: Paulo Alcantara (SUSE) <pc@cjr.nz>
Date:   Mon Jan 13 17:46:59 2020 -0300

    cifs: Fix memory allocation in __smb2_handle_cancelled_cmd()

    commit 0a5a98863c9debc02387b3d23c46d187756f5e2b upstream.

    __smb2_handle_cancelled_cmd() is called under a spin lock held in
    cifs_mid_q_entry_release(), so make its memory allocation GFP_ATOMIC.

    This issue was observed when running xfstests generic/028:

    [ 1722.589204] CIFS VFS: \\192.168.30.26 Cancelling wait for mid 72064 cmd: 5
    [ 1722.590687] CIFS VFS: \\192.168.30.26 Cancelling wait for mid 72065 cmd: 17
    [ 1722.593529] CIFS VFS: \\192.168.30.26 Cancelling wait for mid 72066 cmd: 6
    [ 1723.039014] BUG: sleeping function called from invalid context at mm/slab.h:565
    [ 1723.040710] in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 30877, name: cifsd
    [ 1723.045098] CPU: 3 PID: 30877 Comm: cifsd Not tainted 5.5.0-rc4+ #313
    [ 1723.046256] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.12.0-59-
gc9ba527-rebuilt.opensuse.org 04/01/2014
    [ 1723.048221] Call Trace:
    [ 1723.048689]  dump_stack+0x97/0xe0
    [ 1723.049268]  ___might_sleep.cold+0xd1/0xe1
    [ 1723.050069]  kmem_cache_alloc_trace+0x204/0x2b0
    [ 1723.051051]  __smb2_handle_cancelled_cmd+0x40/0x140 [cifs]
    [ 1723.052137]  smb2_handle_cancelled_mid+0xf6/0x120 [cifs]
    [ 1723.053247]  cifs_mid_q_entry_release+0x44d/0x630 [cifs]
    [ 1723.054351]  ? cifs_reconnect+0x26a/0x1620 [cifs]
    [ 1723.055325]  cifs_demultiplex_thread+0xad4/0x14a0 [cifs]
    [ 1723.056458]  ? cifs_handle_standard+0x2c0/0x2c0 [cifs]
    [ 1723.057365]  ? kvm_sched_clock_read+0x14/0x30
    [ 1723.058197]  ? sched_clock+0x5/0x10
    [ 1723.058838]  ? sched_clock_cpu+0x18/0x110
    [ 1723.059629]  ? lockdep_hardirqs_on+0x17d/0x250
    [ 1723.060456]  kthread+0x1ab/0x200
    [ 1723.061149]  ? cifs_handle_standard+0x2c0/0x2c0 [cifs]
    [ 1723.062078]  ? kthread_create_on_node+0xd0/0xd0
    [ 1723.062897]  ret_from_fork+0x3a/0x50

    Signed-off-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Fixes: 9150c3adbf24 ("CIFS: Close open handle after interrupted close")
    Cc: Stable <stable@vger.kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Reviewed-by: Pavel Shilovsky <pshilov@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b396ec7241edcf75e12900de00f8e8998ed53a1f
Author: Ronnie Sahlberg <lsahlber@redhat.com>
Date:   Wed Jan 8 13:08:07 2020 +1000

    cifs: set correct max-buffer-size for smb2_ioctl_init()

    commit 731b82bb1750a906c1e7f070aedf5505995ebea7 upstream.

    Fix two places where we need to adjust down the max response size for
    ioctl when it is used together with compounding.

    Signed-off-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Reviewed-by: Pavel Shilovsky <pshilov@microsoft.com>
    CC: Stable <stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit d65b067c25465df7468342957dfec3d4788a1807
Author: Vincent Whitchurch <vincent.whitchurch@axis.com>
Date:   Thu Jan 23 17:09:06 2020 +0100

    CIFS: Fix task struct use-after-free on reconnect

    commit f1f27ad74557e39f67a8331a808b860f89254f2d upstream.

    The task which created the MID may be gone by the time cifsd attempts to
    call the callbacks on MIDs from cifs_reconnect().

    This leads to a use-after-free of the task struct in cifs_wake_up_task:

     ==================================================================
     BUG: KASAN: use-after-free in __lock_acquire+0x31a0/0x3270
     Read of size 8 at addr ffff8880103e3a68 by task cifsd/630

     CPU: 0 PID: 630 Comm: cifsd Not tainted 5.5.0-rc6+ #119
     Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.10.2-1 04/01/2014
     Call Trace:
      dump_stack+0x8e/0xcb
      print_address_description.constprop.5+0x1d3/0x3c0
      ? __lock_acquire+0x31a0/0x3270
      __kasan_report+0x152/0x1aa
      ? __lock_acquire+0x31a0/0x3270
      ? __lock_acquire+0x31a0/0x3270
      kasan_report+0xe/0x20
      __lock_acquire+0x31a0/0x3270
      ? __wake_up_common+0x1dc/0x630
      ? _raw_spin_unlock_irqrestore+0x4c/0x60
      ? mark_held_locks+0xf0/0xf0
      ? _raw_spin_unlock_irqrestore+0x39/0x60
      ? __wake_up_common_lock+0xd5/0x130
      ? __wake_up_common+0x630/0x630
      lock_acquire+0x13f/0x330
      ? try_to_wake_up+0xa3/0x19e0
      _raw_spin_lock_irqsave+0x38/0x50
      ? try_to_wake_up+0xa3/0x19e0
      try_to_wake_up+0xa3/0x19e0
      ? cifs_compound_callback+0x178/0x210
      ? set_cpus_allowed_ptr+0x10/0x10
      cifs_reconnect+0xa1c/0x15d0
      ? generic_ip_connect+0x1860/0x1860
      ? rwlock_bug.part.0+0x90/0x90
      cifs_readv_from_socket+0x479/0x690
      cifs_read_from_socket+0x9d/0xe0
      ? cifs_readv_from_socket+0x690/0x690
      ? mempool_resize+0x690/0x690
      ? rwlock_bug.part.0+0x90/0x90
      ? memset+0x1f/0x40
      ? allocate_buffers+0xff/0x340
      cifs_demultiplex_thread+0x388/0x2a50
      ? cifs_handle_standard+0x610/0x610
      ? rcu_read_lock_held_common+0x120/0x120
      ? mark_lock+0x11b/0xc00
      ? __lock_acquire+0x14ed/0x3270
      ? __kthread_parkme+0x78/0x100
      ? lockdep_hardirqs_on+0x3e8/0x560
      ? lock_downgrade+0x6a0/0x6a0
      ? lockdep_hardirqs_on+0x3e8/0x560
      ? _raw_spin_unlock_irqrestore+0x39/0x60
      ? cifs_handle_standard+0x610/0x610
      kthread+0x2bb/0x3a0
      ? kthread_create_worker_on_cpu+0xc0/0xc0
```

```
        ret_from_fork+0x3a/0x50

       Allocated by task 649:
        save_stack+0x19/0x70
        __kasan_kmalloc.constprop.5+0xa6/0xf0
        kmem_cache_alloc+0x107/0x320
        copy_process+0x17bc/0x5370
        _do_fork+0x103/0xbf0
        __x64_sys_clone+0x168/0x1e0
        do_syscall_64+0x9b/0xec0
        entry_SYSCALL_64_after_hwframe+0x49/0xbe

       Freed by task 0:
        save_stack+0x19/0x70
        __kasan_slab_free+0x11d/0x160
        kmem_cache_free+0xb5/0x3d0
        rcu_core+0x52f/0x1230
        __do_softirq+0x24d/0x962

       The buggy address belongs to the object at ffff8880103e32c0
        which belongs to the cache task_struct of size 6016
       The buggy address is located 1960 bytes inside of
        6016-byte region [ffff8880103e32c0, ffff8880103e4a40)
       The buggy address belongs to the page:
       page:ffffea000040f800 refcount:1 mapcount:0 mapping:ffff8880108da5c0
       index:0xffff8880103e4c00 compound_mapcount: 0
       raw: 4000000000010200 ffffea00001f2208 ffffea00001e3408 ffff8880108da5c0
       raw: ffff8880103e4c00 0000000000050003 00000001ffffffff 0000000000000000
       page dumped because: kasan: bad access detected

       Memory state around the buggy address:
        ffff8880103e3900: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
        ffff8880103e3980: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
       >ffff8880103e3a00: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
                                                        ^
        ffff8880103e3a80: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
        ffff8880103e3b00: fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb fb
       ==================================================================

     This can be reliably reproduced by adding the below delay to
     cifs_reconnect(), running find(1) on the mount, restarting the samba
     server while find is running, and killing find during the delay:

             spin_unlock(&GlobalMid_Lock);
             mutex_unlock(&server->srv_mutex);

      +      msleep(10000);
      +
             cifs_dbg(FYI, "%s: issuing mid callbacks\n", __func__);
             list_for_each_safe(tmp, tmp2, &retry_list) {
                     mid_entry = list_entry(tmp, struct mid_q_entry, qhead);

     Fix this by holding a reference to the task struct until the MID is
     freed.

     Signed-off-by: Vincent Whitchurch <vincent.whitchurch@axis.com>
     Signed-off-by: Steve French <stfrench@microsoft.com>
     CC: Stable <stable@vger.kernel.org>
     Reviewed-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
     Reviewed-by: Pavel Shilovsky <pshilov@microsoft.com>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 4f753389bcdc4a5feb753019b0d68ef0922601fb
  Author: Eric Biggers <ebiggers@google.com>
  Date:   Mon Dec 30 21:19:31 2019 -0600
```

```
    crypto: chelsio - fix writing tfm flags to wrong place

    commit bd56cea012fc2d6381e8cd3209510ce09f9de8c9 upstream.

    The chelsio crypto driver is casting 'struct crypto_aead' directly to
    'struct crypto_tfm', which is incorrect because the crypto_tfm isn't the
    first field of 'struct crypto_aead'.  Consequently, the calls to
    crypto_tfm_set_flags() are modifying some other field in the struct.

    Also, the driver is setting CRYPTO_TFM_RES_BAD_KEY_LEN in
    ->setauthsize(), not just in ->setkey().  This is incorrect since this
    flag is for bad key lengths, not for bad authentication tag lengths.

    Fix these bugs by removing the broken crypto_tfm_set_flags() calls from
    ->setauthsize() and by fixing them in ->setkey().

    Fixes: 324429d74127 ("chcr: Support for Chelsio's Crypto Hardware")
    Cc: <stable@vger.kernel.org> # v4.9+
    Cc: Atul Gupta <atul.gupta@chelsio.com>
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9ba7636fdb422594adde839c23db8d764add2343
Author: Guenter Roeck <linux@roeck-us.net>
Date:   Wed Nov 27 12:24:53 2019 -0800

    driver core: Fix test_async_driver_probe if NUMA is disabled

    commit 264d25275a46fce5da501874fa48a2ae5ec571c8 upstream.

    Since commit 57ea974fb871 ("driver core: Rewrite test_async_driver_probe
    to cover serialization and NUMA affinity"), running the test with NUMA
    disabled results in warning messages similar to the following.

    test_async_driver test_async_driver.12: NUMA node mismatch -1 != 0

    If CONFIG_NUMA=n, dev_to_node(dev) returns -1, and numa_node_id()
    returns 0. Both are widely used, so it appears risky to change return
    values. Augment the check with IS_ENABLED(CONFIG_NUMA) instead
    to fix the problem.

    Cc: Alexander Duyck <alexander.h.duyck@linux.intel.com>
    Fixes: 57ea974fb871 ("driver core: Rewrite test_async_driver_probe to cover
serialization and NUMA affinity")
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Cc: stable <stable@vger.kernel.org>
    Acked-by: Alexander Duyck <alexander.h.duyck@linux.intel.com>
    Link: https://lore.kernel.org/r/20191127202453.28087-1-linux@roeck-us.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3811bd2b93afffed5fe99352fcf538c5a3e9e4ca
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Tue Dec 17 19:10:38 2019 +0200

    iio: st_gyro: Correct data for LSM9DS0 gyro

    commit e825070f697abddf3b9b0a675ed0ff1884114818 upstream.

    The commit 41c128cb25ce ("iio: st_gyro: Add lsm9ds0-gyro support")
    assumes that gyro in LSM9DS0 is the same as others with 0xd4 WAI ID,
    but datasheet tells slight different story, i.e. the first scale factor
    for the chip is 245 dps, and not 250 dps.
```

```
    Correct this by introducing a separate settings for LSM9DS0.

    Fixes: 41c128cb25ce ("iio: st_gyro: Add lsm9ds0-gyro support")
    Depends-on: 45a4e4220bf4 ("iio: gyro: st_gyro: fix L3GD20H support")
    Cc: Leonard Crestez <leonard.crestez@nxp.com>
    Cc: Lorenzo Bianconi <lorenzo.bianconi83@gmail.com>
    Cc: <Stable@vger.kernel.org>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a6d19450c3145dd106fc88edc2dea7635856ddf3
Author: Olivier Moysan <olivier.moysan@st.com>
Date:   Wed Nov 27 14:07:29 2019 +0100

    iio: adc: stm32-dfsdm: fix single conversion

    commit dc26935fb60e8da8d59655dd2ec0de47b20d7d8f upstream.

    Apply data formatting to single conversion,
    as this is already done in continuous and trigger modes.

    Fixes: 102afde62937 ("iio: adc: stm32-dfsdm: manage data resolution in trigger mode")

    Signed-off-by: Olivier Moysan <olivier.moysan@st.com>
    Cc: <Stable@vger.kernel.org>
    Acked-by: Fabrice Gasnier <fabrice.gasnier@st.com>
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 29795e4f6a366f5fd16f181ad77369b1f8df0d52
Author: Tomas Winkler <tomas.winkler@intel.com>
Date:   Sun Jan 19 11:42:29 2020 +0200

    mei: me: add comet point (lake) H device ids

    commit 559e575a8946a6561dfe8880de341d4ef78d5994 upstream.

    Add Comet Point device IDs for Comet Lake H platforms.

    Cc: <stable@vger.kernel.org>
    Signed-off-by: Tomas Winkler <tomas.winkler@intel.com>
    Link: https://lore.kernel.org/r/20200119094229.20116-1-tomas.winkler@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9272a0797833eec27403eaa2ccfcf3a4c9283607
Author: Tomas Winkler <tomas.winkler@intel.com>
Date:   Thu Dec 12 10:41:03 2019 +0200

    mei: hdcp: bind only with i915 on the same PCH

    commit 1e8d19d9b0dfcf11b61bac627203a290577e807a upstream.

    The mei device and i915 must reside on the same
    PCH in order for HDCP to work. Make the component
    matching function enforce this requirement.

                        hdcp
                         |
       i915              mei
        |                 |
        +----= PCH =----+

    Cc: <stable@vger.kernel.org> v5.0+
    Cc: Ramalingam C <ramalingam.c@intel.com>
```

```
    Signed-off-by: Tomas Winkler <tomas.winkler@intel.com>
    Reviewed-by: Alexander Usyskin <alexander.usyskin@intel.com>
    Link: https://lore.kernel.org/r/20191212084103.2893-1-tomas.winkler@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7d46883724e04494a4a2df539655fecfd546e827
Author: Martin Fuzzey <martin.fuzzey@flowbird.group>
Date:   Fri Jan 10 16:44:01 2020 +0100

    binder: fix log spam for existing debugfs file creation.

    commit eb143f8756e77c8fcfc4d574922ae9efd3a43ca9 upstream.

    Since commit 43e23b6c0b01 ("debugfs: log errors when something goes wrong")
    debugfs logs attempts to create existing files.

    However binder attempts to create multiple debugfs files with
    the same name when a single PID has multiple contexts, this leads
    to log spamming during an Android boot (17 such messages during
    boot on my system).

    Fix this by checking if we already know the PID and only create
    the debugfs entry for the first context per PID.

    Do the same thing for binderfs for symmetry.

    Signed-off-by: Martin Fuzzey <martin.fuzzey@flowbird.group>
    Acked-by: Todd Kjos <tkjos@google.com>
    Fixes: 43e23b6c0b01 ("debugfs: log errors when something goes wrong")
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/1578671054-5982-1-git-send-email-
martin.fuzzey@flowbird.group
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3e832cf807e8abd40bf38eac04420951728ff8ff
Author: Lubomir Rintel <lkundrak@v3.sk>
Date:   Mon Nov 18 12:54:31 2019 +0100

    component: do not dereference opaque pointer in debugfs

    commit ef9ffc1e5f1ac73ecd2fb3b70db2a3b2472ff2f7 upstream.

    The match data does not have to be a struct device pointer, and indeed
    very often is not. Attempt to treat it as such easily results in a
    crash.

    For the components that are not registered, we don't know which device
    is missing. Once it it is there, we can use the struct component to get
    the device and whether it's bound or not.

    Fixes: 59e73854b5fd ('component: add debugfs support')
    Signed-off-by: Lubomir Rintel <lkundrak@v3.sk>
    Cc: stable <stable@vger.kernel.org>
    Cc: Arnaud Pouliquen <arnaud.pouliquen@st.com>
    Link: https://lore.kernel.org/r/20191118115431.63626-1-lkundrak@v3.sk
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6826af9a5c3f388b53e9f92d8a253f6213341273
Author: Eric Snowberg <eric.snowberg@oracle.com>
Date:   Sat Dec 7 11:16:03 2019 -0500

    debugfs: Return -EPERM when locked down

    commit a37f4958f7b63d2b3cd17a76151fdfc29ce1da5f upstream.
```

```
    When lockdown is enabled, debugfs_is_locked_down returns 1. It will then
    trigger the following:

    WARNING: CPU: 48 PID: 3747
    CPU: 48 PID: 3743 Comm: bash Not tainted 5.4.0-1946.x86_64 #1
    Hardware name: Oracle Corporation ORACLE SERVER X7-2/ASM, MB, X7-2, BIOS 41060400
 05/20/2019
    RIP: 0010:do_dentry_open+0x343/0x3a0
    Code: 00 40 08 00 45 31 ff 48 c7 43 28 40 5b e7 89 e9 02 ff ff ff 48 8b 53 28 4c 8b 72
 70 4d 85 f6 0f 84 10 fe ff ff e9 f5 fd ff ff <0f> 0b 41 bf ea ff ff ff e9 3b ff ff ff 41 bf
 e6 ff ff ff e9 b4 fe
    RSP: 0018:ffffb8740dde7ca0 EFLAGS: 00010202
    RAX: ffffffff89e88a40 RBX: ffff928c8e6b6f00 RCX: 0000000000000000
    RDX: 0000000000000000 RSI: ffff928dbfd97778 RDI: ffff9285cff685c0
    RBP: ffffb8740dde7cc8 R08: 0000000000000821 R09: 0000000000000030
    R10: 0000000000000057 R11: ffffb8740dde7a98 R12: ffff926ec781c900
    R13: ffff928c8e6b6f10 R14: ffffffff8936e190 R15: 0000000000000001
    FS:  00007f45f6777740(0000) GS:ffff928dbfd80000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00007fff95e0d5d8 CR3: 0000001ece562006 CR4: 00000000007606e0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    PKRU: 55555554
    Call Trace:
     vfs_open+0x2d/0x30
     path_openat+0x2d4/0x1680
     ? tty_mode_ioctl+0x298/0x4c0
     do_filp_open+0x93/0x100
     ? strncpy_from_user+0x57/0x1b0
     ? __alloc_fd+0x46/0x150
     do_sys_open+0x182/0x230
     __x64_sys_openat+0x20/0x30
     do_syscall_64+0x60/0x1b0
     entry_SYSCALL_64_after_hwframe+0x170/0x1d5
    RIP: 0033:0x7f45f5e5ce02
    Code: 25 00 00 41 00 3d 00 00 41 00 74 4c 48 8d 05 25 59 2d 00 8b 00 85 c0 75 6d 89 f2
 b8 01 01 00 00 48 89 fe bf 9c ff ff ff 0f 05 <48> 3d 00 f0 ff ff 0f 87 a2 00 00 00 48 8b 4c
 24 28 64 48 33 0c 25
    RSP: 002b:00007fff95e0d2e0 EFLAGS: 00000246 ORIG_RAX: 0000000000000101
    RAX: ffffffffffffffda RBX: 0000561178c069b0 RCX: 00007f45f5e5ce02
    RDX: 0000000000000241 RSI: 0000561178c08800 RDI: 00000000ffffff9c
    RBP: 00007fff95e0d3e0 R08: 0000000000000020 R09: 0000000000000005
    R10: 00000000000001b6 R11: 0000000000000246 R12: 0000000000000000
    R13: 0000000000000003 R14: 0000000000000001 R15: 0000561178c08800

    Change the return type to int and return -EPERM when lockdown is enabled
    to remove the warning above. Also rename debugfs_is_locked_down to
    debugfs_locked_down to make it sound less like it returns a boolean.

    Fixes: 5496197f9b08 ("debugfs: Restrict debugfs when the kernel is locked down")
    Signed-off-by: Eric Snowberg <eric.snowberg@oracle.com>
    Reviewed-by: Matthew Wilcox (Oracle) <willy@infradead.org>
    Cc: stable <stable@vger.kernel.org>
    Acked-by: James Morris <jamorris@linux.microsoft.com>
    Link: https://lore.kernel.org/r/20191207161603.35907-1-eric.snowberg@oracle.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 3b16600e1fd49fbb9ad2a7cdc904807b07403faa
 Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
 Date:   Tue Jan 21 08:17:02 2020 +0100

    serial: imx: fix a race condition in receive path

    commit 101aa46bd221b768dfff8ef3745173fc8dbb85ee upstream.
```

```
    The main irq handler function starts by first masking disabled
    interrupts in the status register values to ensure to only handle
    enabled interrupts. This is important as when the RX path in the
    hardware is disabled reading the RX fifo results in an external abort.

    This checking must be done under the port lock, otherwise the following
    can happen:

         CPU1                             | CPU2
                                          |
         irq triggers as there are chars  |
         in the RX fifo                   |
                                          | grab port lock
         imx_uart_int finds RRDY enabled  |
         and calls imx_uart_rxint which   |
         has to wait for port lock        |
                                          | disable RX (e.g. because we're
                                          | using RS485 with !RX_DURING_TX)
                                          |
                                          | release port lock
         read from RX fifo with RX        |
         disabled => exception            |

    So take the port lock only once in imx_uart_int() instead of in the
    functions called from there.

    Reported-by: Andre Renaud <arenaud@designa-electronics.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Link: https://lore.kernel.org/r/20200121071702.20150-1-u.kleine-koenig@pengutronix.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f54b5d4a1abc9e5047320f8dda2e3f4df0fd9757
Author: Lukas Wunner <lukas@wunner.de>
Date:   Thu Jan 16 13:14:01 2020 +0100

    serial: 8250_bcm2835aux: Fix line mismatch on driver unbind

    commit dc76697d7e933d5e299116f219c890568785ea15 upstream.

    Unbinding the bcm2835aux UART driver raises the following error if the
    maximum number of 8250 UARTs is set to 1 (via the 8250.nr_uarts module
    parameter or CONFIG_SERIAL_8250_RUNTIME_UARTS):

    (NULL device *): Removing wrong port: a6f80333 != fa20408b

    That's because bcm2835aux_serial_probe() retrieves UART line number 1
    from the devicetree and stores it in data->uart.port.line, while
    serial8250_register_8250_port() instead uses UART line number 0,
    which is stored in data->line.

    On driver unbind, bcm2835aux_serial_remove() uses data->uart.port.line,
    which contains the wrong number.  Fix it.

    The issue does not occur if the maximum number of 8250 UARTs is >= 2.

    Fixes: bdc5f3009580 ("serial: bcm2835: add driver for bcm2835-aux-uart")
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Cc: stable@vger.kernel.org # v4.6+
    Cc: Martin Sperl <kernel@martin.sperl.org>
    Reviewed-by: Nicolas Saenz Julienne <nsaenzjulienne@suse.de>
    Tested-by: Nicolas Saenz Julienne <nsaenzjulienne@suse.de>
    Link:
https://lore.kernel.org/r/912ccf553c5258135c6d7e8f404a101ef320f0f4.1579175223.git.lukas@wunn
er.de
```

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

commit 0710da067387eea41612f4a821b1d92d6c7ad88c
Author: Malcolm Priestley <tvboxspy@gmail.com>
Date:   Wed Jan 8 21:41:36 2020 +0000

    staging: vt6656: Fix false Tx excessive retries reporting.

    commit 9dd631fa99dc0a0dfbd191173bf355ba30ea786a upstream.

    The driver reporting  IEEE80211_TX_STAT_ACK is not being handled
    correctly. The driver should only report on TSR_TMO flag is not
    set indicating no transmission errors and when not IEEE80211_TX_CTL_NO_ACK
    is being requested.

    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Malcolm Priestley <tvboxspy@gmail.com>
    Link: https://lore.kernel.org/r/340f1f7f-c310-dca5-476f-abc059b9cd97@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a56c0f7d5f1506481d359ed098c6a61f7fe5f80b
Author: Malcolm Priestley <tvboxspy@gmail.com>
Date:   Wed Jan 8 21:41:20 2020 +0000

    staging: vt6656: use NULLFUCTION stack on mac80211

    commit d579c43c82f093e63639151625b2139166c730fd upstream.

    It appears that the drivers does not go into power save correctly the
    NULL data packets are not being transmitted because it not enabled
    in mac80211.

    The driver needs to capture ieee80211_is_nullfunc headers and
    copy the duration_id to it's own duration data header.

    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Malcolm Priestley <tvboxspy@gmail.com>
    Link: https://lore.kernel.org/r/610971ae-555b-a6c3-61b3-444a0c1e35b4@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1a80cc38db42b0bd8ce474116f5bf3d11a357e6f
Author: Malcolm Priestley <tvboxspy@gmail.com>
Date:   Wed Jan 8 21:40:58 2020 +0000

    staging: vt6656: correct packet types for CTS protect, mode.

    commit d971fdd3412f8342747778fb59b8803720ed82b1 upstream.

    It appears that the driver still transmits in CTS protect mode even
    though it is not enabled in mac80211.

    That is both packet types PK_TYPE_11GA and PK_TYPE_11GB both use CTS protect.
    The only difference between them GA does not use B rates.

    Find if only B rate in GB or GA in protect mode otherwise transmit packets
    as PK_TYPE_11A.

    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Malcolm Priestley <tvboxspy@gmail.com>
    Link: https://lore.kernel.org/r/9c1323ff-dbb3-0eaa-43e1-9453f7390dc0@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3a5bff77f48c869623724a2dfbba0bdeb1d3e969
Author: Colin Ian King <colin.king@canonical.com>
Date:   Tue Jan 14 18:16:04 2020 +0000
```

```
    staging: wlan-ng: ensure error return is actually returned

    commit 4cc41cbce536876678b35e03c4a8a7bb72c78fa9 upstream.

    Currently when the call to prism2sta_ifst fails a netdev_err error
    is reported, error return variable result is set to -1 but the
    function always returns 0 for success.  Fix this by returning
    the error value in variable result rather than 0.

    Addresses-Coverity: ("Unused value")
    Fixes: 00b3ed168508 ("Staging: add wlan-ng prism2 usb driver")
    Signed-off-by: Colin Ian King <colin.king@canonical.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20200114181604.390235-1-colin.king@canonical.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b5e5d81230ec6a24b3ce452fc41d8260292c686a
Author: Andrey Shvetsov <andrey.shvetsov@k2l.de>
Date:   Thu Jan 16 18:22:39 2020 +0100

    staging: most: net: fix buffer overflow

    commit 4d1356ac12f4d5180d0df345d85ff0ee42b89c72 upstream.

    If the length of the socket buffer is 0xFFFFFFFF (max size for an
    unsigned int), then payload_len becomes 0xFFFFFFF1 after subtracting 14
    (ETH_HLEN).  Then, mdp_len is set to payload_len + 16 (MDP_HDR_LEN)
    which overflows and results in a value of 2.  These values for
    payload_len and mdp_len will pass current buffer size checks.

    This patch checks if derived from skb->len sum may overflow.

    The check is based on the following idea:

    For any `unsigned V1, V2` and derived `unsigned SUM = V1 + V2`,
    `V1 + V2` overflows iif `SUM < V1`.

    Reported-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Andrey Shvetsov <andrey.shvetsov@k2l.de>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20200116172238.6046-1-andrey.shvetsov@microchip.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9ba5957e034dea637224bea219b0baf5b354a819
Author: Thomas Hebb <tommyhebb@gmail.com>
Date:   Mon Jan 20 06:09:06 2020 -0800

    usb: typec: fusb302: fix "op-sink-microwatt" default that was in mW

    commit eb7a3bb8c955b3694e0e0998413ce1563c02f90c upstream.

    commit 8f6244055bd3 ("usb: typec: fusb302: Always provide fwnode for the
    port") didn't convert this value from mW to uW when migrating to a new
    specification format like it should have.

    Fixes: 8f6244055bd3 ("usb: typec: fusb302: Always provide fwnode for the port")
    Cc: stable@vger.kernel.org
    Signed-off-by: Thomas Hebb <tommyhebb@gmail.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link:
https://lore.kernel.org/r/0da564559af75ec829c6c7e3aa4024f857c91bee.1579529334.git.tommyhebb@
gmail.com
```

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

commit 9bfa308e95b1593c109ce860ae749fe5c2ccecb6
Author: Thomas Hebb <tommyhebb@gmail.com>
Date:   Mon Jan 20 06:09:05 2020 -0800

    usb: typec: wcove: fix "op-sink-microwatt" default that was in mW

    commit 0e64350bf4668d0fbbfec66fd8e637b971b4e976 upstream.

    commit 4c912bff46cc ("usb: typec: wcove: Provide fwnode for the port")
    didn't convert this value from mW to uW when migrating to a new
    specification format like it should have.

    Fixes: 4c912bff46cc ("usb: typec: wcove: Provide fwnode for the port")
    Cc: stable@vger.kernel.org
    Signed-off-by: Thomas Hebb <tommyhebb@gmail.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link:
https://lore.kernel.org/r/d8be32512efd31995ad7d65b27df9d443131b07c.1579529334.git.tommyhebb@
gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 445839bc7fa3305bdf57dce07ebf8216c35d25a4
Author: Bin Liu <b-liu@ti.com>
Date:   Wed Dec 11 10:10:03 2019 -0600

    usb: dwc3: turn off VBUS when leaving host mode

    commit 09ed259fac621634d51cd986aa8d65f035662658 upstream.

    VBUS should be turned off when leaving the host mode.
    Set GCTL_PRTCAP to device mode in teardown to de-assert DRVVBUS pin to
    turn off VBUS power.

    Fixes: 5f94adfeed97 ("usb: dwc3: core: refactor mode initialization to its own
function")
    Cc: stable@vger.kernel.org
    Signed-off-by: Bin Liu <b-liu@ti.com>
    Signed-off-by: Felipe Balbi <balbi@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1fe175772b993c1334bcfc030a1486656778507d
Author: Johan Hovold <johan@kernel.org>
Date:   Wed Jan 22 11:15:28 2020 +0100

    USB: serial: ir-usb: fix IrLAP framing

    commit 38c0d5bdf4973f9f5a888166e9d3e9ed0d32057a upstream.

    Commit f4a4cbb2047e ("USB: ir-usb: reimplement using generic framework")
    switched to using the generic write implementation which may combine
    multiple write requests into larger transfers. This can break the IrLAP
    protocol where end-of-frame is determined using the USB short packet
    mechanism, for example, if multiple frames are sent in rapid succession.

    Fixes: f4a4cbb2047e ("USB: ir-usb: reimplement using generic framework")
    Cc: stable <stable@vger.kernel.org>     # 2.6.35
    Reviewed-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c9b7dcee64c1cd414b5e1c6d3b561753956c7d6c
Author: Johan Hovold <johan@kernel.org>
```

```
Date:   Wed Jan 22 11:15:27 2020 +0100

    USB: serial: ir-usb: fix link-speed handling

    commit 17a0184ca17e288decdca8b2841531e34d49285f upstream.

    Commit e0d795e4f36c ("usb: irda: cleanup on ir-usb module") added a USB
    IrDA header with common defines, but mistakingly switched to using the
    class-descriptor baud-rate bitmask values for the outbound header.

    This broke link-speed handling for rates above 9600 baud, but a device
    would also be able to operate at the default 9600 baud until a
    link-speed request was issued (e.g. using the TCGETS ioctl).

    Fixes: e0d795e4f36c ("usb: irda: cleanup on ir-usb module")
    Cc: stable <stable@vger.kernel.org>     # 2.6.27
    Cc: Felipe Balbi <balbi@kernel.org>
    Reviewed-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 85eec8cf9d2232db69416d115ea91baa7a40f6da
Author: Johan Hovold <johan@kernel.org>
Date:   Wed Jan 22 11:15:26 2020 +0100

    USB: serial: ir-usb: add missing endpoint sanity check

    commit 2988a8ae7476fe9535ab620320790d1714bdad1d upstream.

    Add missing endpoint sanity check to avoid dereferencing a NULL-pointer
    on open() in case a device lacks a bulk-out endpoint.

    Note that prior to commit f4a4cbb2047e ("USB: ir-usb: reimplement using
    generic framework") the oops would instead happen on open() if the
    device lacked a bulk-in endpoint and on write() if it lacked a bulk-out
    endpoint.

    Fixes: f4a4cbb2047e ("USB: ir-usb: reimplement using generic framework")
    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Cc: stable <stable@vger.kernel.org>
    Reviewed-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 380a352cd3e9bb3a912d0dc68d996033ec7a843d
Author: Peter Robinson <pbrobinson@gmail.com>
Date:   Mon Jan 20 14:19:10 2020 +0000

    usb: host: xhci-tegra: set MODULE_FIRMWARE for tegra186

    commit e1f236efd9c579a29d7df75aa052127d0d975267 upstream.

    Set the MODULE_FIRMWARE for tegra186, it's registered for 124/210 and
    ensures the firmware is available at the appropriate time such as in
    the initrd, else if the firmware is unavailable the driver fails with
    the following errors:

    tegra-xusb 3530000.usb: Direct firmware load for nvidia/tegra186/xusb.bin failed with
error -2
    tegra-xusb 3530000.usb: failed to request firmware: -2
    tegra-xusb 3530000.usb: failed to load firmware: -2
    tegra-xusb: probe of 3530000.usb failed with error -2

    Fixes: 5f9be5f3f899 ("usb: host: xhci-tegra: Add Tegra186 XUSB support")
    Signed-off-by: Peter Robinson <pbrobinson@gmail.com>
```

```
    Acked-by: Thierry Reding <treding@nvidia.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20200120141910.116097-1-pbrobinson@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f841bbcf8130f34b253d3e8c84c9073d00779440
Author: Heikki Krogerus <heikki.krogerus@linux.intel.com>
Date:   Fri Jan 17 12:30:33 2020 +0300

    usb: dwc3: pci: add ID for the Intel Comet Lake -V variant

    commit f5ae8869095552e3396ee3e404f9586cc6a828f0 upstream.

    There is one more Comet Lake PCH variant, CML-V, that has
    its own PCI ID.

    Signed-off-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20200117093033.48616-1-heikki.krogerus@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 566264c8b6e3b85dfa9febcaa5e1920f0ee6c4ba
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:25 2019 +0100

    rsi_91x_usb: fix interface sanity check

    commit 3139b180906af43bc09bd3373fc2338a8271d9d9 upstream.

    Make sure to use the current alternate setting when verifying the
    interface descriptors to avoid binding to an invalid interface.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: dad0d04fa7ba ("rsi: Add RS9113 wireless driver")
    Cc: stable <stable@vger.kernel.org>     # 3.15
    Cc: Fariya Fatima <fariyaf@gmail.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c48a89147160f7d7ee9682242bcfc6f337f1fa0d
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Dec 10 12:44:23 2019 +0100

    orinoco_usb: fix interface sanity check

    commit b73e05aa543cf8db4f4927e36952360d71291d41 upstream.

    Make sure to use the current alternate setting when verifying the
    interface descriptors to avoid binding to an invalid interface.

    Failing to do so could cause the driver to misbehave or trigger a WARN()
    in usb_submit_urb() that kernels with panic_on_warn set would choke on.

    Fixes: 9afac70a7305 ("orinoco: add orinoco_usb driver")
    Cc: stable <stable@vger.kernel.org>     # 2.6.35
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 619c458f8f398ab3d54d5a2b47721c996f5da4c7
Author: Johan Hovold <johan@kernel.org>
Date:   Thu Nov 28 19:24:27 2019 +0100
```

```
    Bluetooth: btusb: fix non-atomic allocation in completion handler

    commit 22cc6b7a1dbb58da4afc539d9b7d470b23a25eea upstream.

    USB completion handlers are called in atomic context and must
    specifically not allocate memory using GFP_KERNEL.

    Fixes: a1c49c434e15 ("Bluetooth: btusb: Add protocol support for MediaTek MT7668U USB
devices")
    Cc: stable <stable@vger.kernel.org>     # 5.3
    Cc: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```