```
commit 51dd976781da8c0b47e106ed59a96d7c28972ce4
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Wed Oct 12 09:51:26 2022 +0200

    Linux 5.19.15

    Link: https://lore.kernel.org/r/20221010070333.676316214@linuxfoundation.org
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Link: https://lore.kernel.org/r/20221010191212.200768859@linuxfoundation.org
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Slade Watkins <srw@sladewatkins.net>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Sudip Mukherjee <sudip.mukherjee@codethink.co.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 04bd2a779511e04892b522d41abbdb0e5c072abf
Author: Sasha Levin <sashal@kernel.org>
Date:   Tue Oct 11 08:58:00 2022 -0400

    Revert "clk: ti: Stop using legacy clkctrl names for omap4 and 5"

    This reverts commit 67c830a6de835a36b4e19fe4d968dbaf8dc4e9c6.

    Which was upstream commit 255584b138343d4a28c6d25bd82d04b09460d672.

    Reported as causing boot failures.

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

commit c4635cf3d845a7324c25c52d549b70c8bd7ad4c7
Author: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
Date:   Fri Sep 2 20:23:48 2022 +0900

    Bluetooth: use hdev->workqueue when queuing hdev->{cmd,ncmd}_timer works

    commit deee93d13d385103205879a8a0915036ecd83261 upstream.

    syzbot is reporting attempt to schedule hdev->cmd_work work from system_wq
    WQ into hdev->workqueue WQ which is under draining operation [1], for
    commit c8efcc2589464ac7 ("workqueue: allow chained queueing during
    destruction") does not allow such operation.

    The check introduced by commit 877afadad2dce8aa ("Bluetooth: When HCI work
    queue is drained, only queue chained work") was incomplete.

    Use hdev->workqueue WQ when queuing hdev->{cmd,ncmd}_timer works because
    hci_{cmd,ncmd}_timeout() calls queue_work(hdev->workqueue). Also, protect
    the queuing operation with RCU read lock in order to avoid calling
    queue_delayed_work() after cancel_delayed_work() completed.

    Link: https://syzkaller.appspot.com/bug?extid=243b7d89777f90f7613b [1]
    Reported-by: syzbot <syzbot+243b7d89777f90f7613b@syzkaller.appspotmail.com>
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Fixes: 877afadad2dce8aa ("Bluetooth: When HCI work queue is drained, only queue chained
work")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
```

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

commit 6954cee8826819340594ef614341e4d634af3d4a
Author: Jules Irenge <jbi.octave@gmail.com>
Date:   Wed Sep 7 16:24:20 2022 +0100

    bpf: Fix resetting logic for unreferenced kptrs

    commit 9fad7fe5b29803584c7f17a2abe6c2936fec6828 upstream.

    Sparse reported a warning at bpf_map_free_kptrs()
    "warning: Using plain integer as NULL pointer"
    During the process of fixing this warning, it was discovered that the current
    code erroneously writes to the pointer variable instead of deferencing and
    writing to the actual kptr. Hence, Sparse tool accidentally helped to uncover
    this problem. Fix this by doing WRITE_ONCE(*p, 0) instead of WRITE_ONCE(p, 0).

    Note that the effect of this bug is that unreferenced kptrs will not be cleared
    during check_and_free_fields. It is not a problem if the clearing is not done
    during map_free stage, as there is nothing to free for them.

    Fixes: 14a324f6a67e ("bpf: Wire up freeing of referenced kptr")
    Signed-off-by: Jules Irenge <jbi.octave@gmail.com>
    Link: https://lore.kernel.org/r/Yxi3pJaK6UDjVJSy@playground
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 292b46c49675fcff39547b643d48d0fec27cdbb9
Author: Daniel Golle <daniel@makrotopia.org>
Date:   Fri Sep 30 01:56:53 2022 +0100

    net: ethernet: mtk_eth_soc: fix state in __mtk_foe_entry_clear

    commit ae3ed15da5889263de372ff9df2e83e16acca4cb upstream.

    Setting ib1 state to MTK_FOE_STATE_UNBIND in __mtk_foe_entry_clear
    routine as done by commit 0e80707d94e4c8 ("net: ethernet: mtk_eth_soc:
    fix typo in __mtk_foe_entry_clear") breaks flow offloading, at least
    on older MTK_NETSYS_V1 SoCs, OpenWrt users have confirmed the bug on
    MT7622 and MT7621 systems.
    Felix Fietkau suggested to use MTK_FOE_STATE_INVALID instead which
    works well on both, MTK_NETSYS_V1 and MTK_NETSYS_V2.

    Tested on MT7622 (Linksys E8450) and MT7986 (BananaPi BPI-R3).

    Suggested-by: Felix Fietkau <nbd@nbd.name>
    Fixes: 0e80707d94e4c8 ("net: ethernet: mtk_eth_soc: fix typo in __mtk_foe_entry_clear")
    Fixes: 33fc42de33278b ("net: ethernet: mtk_eth_soc: support creating mac address based
offload entries")
    Signed-off-by: Daniel Golle <daniel@makrotopia.org>
    Link: https://lore.kernel.org/r/YzY+1Yg0FBXcnrtc@makrotopia.org
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9783292253fdc4001268729e1bd08777ba9248c5
Author: Kumar Kartikeya Dwivedi <memxor@gmail.com>
Date:   Wed Sep 21 16:35:50 2022 +0200

    bpf: Gate dynptr API behind CAP_BPF

    commit 8addbfc7b308d591f8a5f2f6bb24d08d9d79dfbb upstream.

    This has been enabled for unprivileged programs for only one kernel
    release, hence the expected annoyances due to this move are low. Users
    using ringbuf can stick to non-dynptr APIs. The actual use cases dynptr
```

```
    is meant to serve may not make sense in unprivileged BPF programs.

    Hence, gate these helpers behind CAP_BPF and limit use to privileged
    BPF programs.

    Fixes: 263ae152e962 ("bpf: Add bpf_dynptr_from_mem for local dynptrs")
    Fixes: bc34dee65a65 ("bpf: Dynptr support for ring buffers")
    Fixes: 13bbbfbea759 ("bpf: Add bpf_dynptr_read and bpf_dynptr_write")
    Fixes: 34d4ef5775f7 ("bpf: Add dynptr data slices")
    Signed-off-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20220921143550.30247-1-memxor@gmail.com
    Acked-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7d3f5d0353b4d30216a6c6eedba962c3279fa635
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu May 19 09:33:28 2022 +0200

    rpmsg: qcom: glink: replace strncpy() with strscpy_pad()

    commit 766279a8f85df32345dbda03b102ca1ee3d5ddea upstream.

    The use of strncpy() is considered deprecated for NUL-terminated
    strings[1]. Replace strncpy() with strscpy_pad(), to keep existing
    pad-behavior of strncpy, similarly to commit 08de420a8014 ("rpmsg:
    glink: Replace strncpy() with strscpy_pad()").  This fixes W=1 warning:

      In function ‘qcom_glink_rx_close’,
        inlined from ‘qcom_glink_work’ at ../drivers/rpmsg/qcom_glink_native.c:1638:4:
      drivers/rpmsg/qcom_glink_native.c:1549:17: warning: ‘strncpy’ specified bound 32
equals destination size [-Wstringop-truncation]
       1549 |                 strncpy(chinfo.name, channel->name, sizeof(chinfo.name));

    [1] https://www.kernel.org/doc/html/latest/process/deprecated.html#strncpy-on-nul-
terminated-strings

    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20220519073330.7187-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Andrew Chernyakov <acherniakov@astralinux.ru>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 165ae42fc1cf3c27a4cd077d4622eec4d5621914
Author: Brian Norris <briannorris@chromium.org>
Date:   Tue Sep 13 18:40:10 2022 -0700

    mmc: core: Terminate infinite loop in SD-UHS voltage switch

    commit e9233917a7e53980664efbc565888163c0a33c3f upstream.

    This loop intends to retry a max of 10 times, with some implicit
    termination based on the SD_{R,}OCR_S18A bit. Unfortunately, the
    termination condition depends on the value reported by the SD card
    (*rocr), which may or may not correctly reflect what we asked it to do.

    Needless to say, it's not wise to rely on the card doing what we expect;
    we should at least terminate the loop regardless. So, check both the
    input and output values, so we ensure we will terminate regardless of
    the SD card behavior.

    Note that SDIO learned a similar retry loop in commit 0797e5f1453b
    ("mmc: core: Fixup signal voltage switch"), but that used the 'ocr'
    result, and so the current pre-terminating condition looks like:
```

```
        rocr & ocr & R4_18V_PRESENT

    (i.e., it doesn't have the same bug.)

    This addresses a number of crash reports seen on ChromeOS that look
    like the following:

        ... // lots of repeated: ...
        <4>[13142.846061] mmc1: Skipping voltage switch
        <4>[13143.406087] mmc1: Skipping voltage switch
        <4>[13143.964724] mmc1: Skipping voltage switch
        <4>[13144.526089] mmc1: Skipping voltage switch
        <4>[13145.086088] mmc1: Skipping voltage switch
        <4>[13145.645941] mmc1: Skipping voltage switch
        <3>[13146.153969] INFO: task halt:30352 blocked for more than 122 seconds.
        ...

    Fixes: f2119df6b764 ("mmc: sd: add support for signal voltage switch procedure")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20220914014010.2076169-1-briannorris@chromium.org
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b2afdaddaa8c33818986cc8e355de566636328b1
Author: ChanWoo Lee <cw9316.lee@samsung.com>
Date:   Wed Jul 6 09:48:40 2022 +0900

    mmc: core: Replace with already defined values for readability

    commit e427266460826bea21b70f9b2bb29decfb2c2620 upstream.

    SD_ROCR_S18A is already defined and is used to check the rocr value, so
    let's replace with already defined values for readability.

    Signed-off-by: ChanWoo Lee <cw9316.lee@samsung.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20220706004840.24812-1-cw9316.lee@samsung.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0b0f40745ab405777e9df8d8a46dd4302f40d45c
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Tue Aug 2 23:25:00 2022 -0500

    gpiolib: acpi: Add a quirk for Asus UM325UAZ

    commit 0ea76c401f9245ac209f1b1ce03a7e1fb9de36e5 upstream.

    Asus UM325UAZ has GPIO 18 programmed as both an interrupt and a wake
    source, but confirmed with internal team on this design this pin is
    floating and shouldn't have been programmed. This causes lots of
    spurious IRQs on the system and horrendous battery life.

    Add a quirk to ignore attempts to program this pin on this system.

    Reported-by: Pavel Krc <reg.krn@pkrc.net>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=216208
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Reviewed-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
```

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

commit aa476c7b637521457b099edb18b2c33d1c645109
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Tue Aug 2 23:24:59 2022 -0500

    gpiolib: acpi: Add support to ignore programming an interrupt

    commit 6b6af7bd5718f4e45a9b930533aec1158387d552 upstream.

    gpiolib-acpi already had support for ignoring a pin for wakeup, but
    if an OEM configures a floating pin as an interrupt source then
    stopping it from being a wakeup won't do much good to stop the
    interrupt storm.

    Add support for a module parameter and quirk infrastructure to
    ignore interrupts as well.

    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a9d2ce58126ca8b7f2a2a8a863c84b857100fe9c
Author: Johan Hovold <johan@kernel.org>
Date:   Tue Sep 13 16:53:12 2022 +0200

    USB: serial: ftdi_sio: fix 300 bps rate for SIO

    commit 7bd7ad3c310cd6766f170927381eea0aa6f46c69 upstream.

    The 300 bps rate of SIO devices has been mapped to 9600 bps since
    2003... Let's fix the regression.

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

commit 9de74019cd44ddddaaca3c677343750ac78b6f73
Author: Tadeusz Struk <tadeusz.struk@linaro.org>
Date:   Mon Sep 19 14:59:57 2022 -0700

    usb: mon: make mmapped memory read only

    commit a659daf63d16aa883be42f3f34ff84235c302198 upstream.

    Syzbot found an issue in usbmon module, where the user space client can
    corrupt the monitor's internal memory, causing the usbmon module to
    crash the kernel with segfault, UAF, etc.

    The reproducer mmaps the /dev/usbmon memory to user space, and
    overwrites it with arbitrary data, which causes all kinds of issues.

    Return an -EPERM error from mon_bin_mmap() if the flag VM_WRTIE is set.
    Also clear VM_MAYWRITE to make it impossible to change it to writable
    later.

    Cc: "Dmitry Vyukov" <dvyukov@google.com>
    Cc: stable <stable@kernel.org>
    Fixes: 6f23ee1fefdc ("USB: add binary API to usbmon")
    Suggested-by: PaX Team <pageexec@freemail.hu>   # for the VM_MAYRITE portion
    Link: https://syzkaller.appspot.com/bug?id=2eb1f35d6525fa4a74d75b4244971e5b1411c95a
    Reported-by: syzbot+23f57c5ae902429285d7@syzkaller.appspotmail.com
    Signed-off-by: Tadeusz Struk <tadeusz.struk@linaro.org>
```

```
    Link: https://lore.kernel.org/r/20220919215957.205681-1-tadeusz.struk@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c2423912d5c93fc25fec1f11699b4406cc07e4f7
Author: Zhang Qilong <zhangqilong3@huawei.com>
Date:   Thu Sep 29 22:30:38 2022 +0800

    i2c: davinci: fix PM disable depth imbalance in davinci_i2c_probe

    [ Upstream commit e2062df704dea47efe16edcaa2316d7b5ecca64f ]

    The pm_runtime_enable will increase power disable depth. Thus a
    pairing decrement is needed on the error handling path to keep
    it balanced according to context.

    Fixes: 17f88151ff190 ("i2c: davinci: Add PM Runtime Support")
    Signed-off-by: Zhang Qilong <zhangqilong3@huawei.com>
    Reviewed-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71c41e04d9d8f89525bd153e80f6b4ac01ee0834
Author: Al Viro <viro@zeniv.linux.org.uk>
Date:   Mon Sep 26 11:59:14 2022 -0400

    don't use __kernel_write() on kmap_local_page()

    [ Upstream commit 06bbaa6dc53cb72040db952053432541acb9adc7 ]

    passing kmap_local_page() result to __kernel_write() is unsafe -
    random ->write_iter() might (and 9p one does) get unhappy when
    passed ITER_KVEC with pointer that came from kmap_local_page().

    Fix by providing a variant of __kernel_write() that takes an iov_iter
    from caller (__kernel_write() becomes a trivial wrapper) and adding
    dump_emit_page() that parallels dump_emit(), except that instead of
    __kernel_write() it uses __kernel_write_iter() with ITER_BVEC source.

    Fixes: 3159ed57792b "fs/coredump: use kmap_local_page()"
    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5dac20360a7b4acc8f418198edee259972cc7c9b
Author: Kan Liang <kan.liang@linux.intel.com>
Date:   Wed Aug 31 07:27:02 2022 -0700

    perf/x86/intel: Fix unchecked MSR access error for Alder Lake N

    [ Upstream commit 24919fdea6f8b31d7cdf32ac291bc5dd0b023878 ]

    For some Alder Lake N machine, the below unchecked MSR access error may be
    triggered.

    [ 0.088017] rcu: Hierarchical SRCU implementation.
    [ 0.088017] unchecked MSR access error: WRMSR to 0x38f (tried to write
    0x0001000f0000003f) at rIP: 0xffffffffb5684de8 (native_write_msr+0x8/0x30)
    [ 0.088017] Call Trace:
    [ 0.088017] <TASK>
    [ 0.088017] __intel_pmu_enable_all.constprop.46+0x4a/0xa0

    The Alder Lake N only has e-cores. The X86_FEATURE_HYBRID_CPU flag is
    not set. The perf cannot retrieve the correct CPU type via
    get_this_hybrid_cpu_type(). The model specific get_hybrid_cpu_type() is
    hardcode to p-core. The wrong CPU type is given to the PMU of the
    Alder Lake N.
```

```
    Since Alder Lake N isn't in fact a hybrid CPU, remove ALDERLAKE_N from
    the rest of {ALDER,RAPTOP}LAKE and create a non-hybrid PMU setup.

    The differences between Gracemont and the previous Tremont are,
    - Number of GP counters
    - Load and store latency Events
    - PEBS event_constraints
    - Instruction Latency support
    - Data source encoding
    - Memory access latency encoding

    Fixes: c2a960f7c574 ("perf/x86: Add new Alder Lake and Raptor Lake support")
    Reported-by: Jianfeng Gao <jianfeng.gao@intel.com>
    Suggested-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lkml.kernel.org/r/20220831142702.153110-1-kan.liang@linux.intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3682460b932761598f2e0254ef6ed25b2064b990
Author: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
Date:   Wed Jul 27 13:02:44 2022 -0400

    drm/amd/display: increase dcn315 pstate change latency

    [ Upstream commit dcc2527df918edfe297c5074ccc1f05eae361ca6 ]

    [Why & How]
    Update after new measurment came in

    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94d49c380b9a8e6da6a05d2997f5fa637e294ae7
Author: Cruise Hung <Cruise.Hung@amd.com>
Date:   Thu Sep 8 22:04:09 2022 +0800

    drm/amd/display: Fix DP MST timeslot issue when fallback happened

    [ Upstream commit 20c6168b3c8aadef7d2853c925d99eb546bd5e1c ]

    [Why]
    When USB4 DP link training failed and fell back to lower link rate,
    the time slot calculation uses the verified_link_cap.
    And the verified_link_cap was not updated to the new one.
    It caused the wrong VC payload time-slot was allocated.

    [How]
    Updated verified_link_cap with the new one from cur_link_settings
    after the LT completes successfully.

    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: Cruise Hung <Cruise.Hung@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fb0b18590c660c30c949da57966a99c4395e188
Author: zhikzhai <zhikai.zhai@amd.com>
```

```
Date:   Fri Aug 26 19:44:50 2022 +0800

    drm/amd/display: skip audio setup when audio stream is enabled

    [ Upstream commit 65fbfb02c2734cacffec5e3f492e1b4f1dabcf98 ]

    [why]
    We have minimal pipe split transition method to avoid pipe
    allocation outage.However, this method will invoke audio setup
    which cause audio output stuck once pipe reallocate.

    [how]
    skip audio setup for pipelines which audio stream has been enabled

    Reviewed-by: Charlene Liu <Charlene.Liu@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: zhikzhai <zhikai.zhai@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c15b295ae2f0c1ddd310c71402b04310f157e6b
Author: Hugo Hu <hugo.hu@amd.com>
Date:   Thu Sep 1 10:08:35 2022 +0800

    drm/amd/display: update gamut remap if plane has changed

    [ Upstream commit 52bb21499cf54fa65b56d97cd0d68579c90207dd ]

    [Why]
    The desktop plane and full-screen game plane may have different
    gamut remap coefficients, if switching between desktop and
    full-screen game without updating the gamut remap will cause
    incorrect color.

    [How]
    Update gamut remap if planes change.

    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: Hugo Hu <hugo.hu@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89d3870658033807941fea5ffa5143cd9526ace9
Author: Michael Strauss <michael.strauss@amd.com>
Date:   Wed Aug 31 15:10:43 2022 -0400

    drm/amd/display: Assume an LTTPR is always present on fixed_vs links

    [ Upstream commit 29956d0fded036a570bd8e7d4ea4b1a1730307d2 ]

    [WHY]
    LTTPRs can in very rare instsances fail to increment DPCD LTTPR count.
    This results in aux-i LTTPR requests to be sent to the wrong DPCD
    address, which causes link training failure.

    [HOW]
    Override internal repeater count if fixed_vs flag is set for a given link

    Reviewed-by: George Shen <George.Shen@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: Michael Strauss <michael.strauss@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
```

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

commit d1137c3c03d83c8705863b1b22dc178e0be93d48
Author: Leo Li <sunpeng.li@amd.com>
Date:   Tue Aug 30 16:38:16 2022 -0400

    drm/amd/display: Fix double cursor on non-video RGB MPO

    [ Upstream commit b261509952bc19d1012cf732f853659be6ebc61e ]

    [Why]

    DC makes use of layer_index (zpos) when picking the HW plane to enable
    HW cursor on. However, some compositors will not attach zpos information
    to each DRM plane. Consequently, in amdgpu, we default layer_index to 0
    and do not update it.

    This causes said DC logic to enable HW cursor on all planes of the same
    layer_index, which manifests as a double cursor issue if one of the
    planes is scaled (and hence scaling the cursor as well).

    [How]

    Use DRM core helpers to calculate a normalized_zpos value for each
    drm_plane_state under each crtc, within the atomic state.

    This helper will first consider existing zpos values, and if
    identical/unset, fallback to plane ID ordering.

    The normalized_zpos is then passed to dc_plane_info during atomic check
    for later use by the cursor logic.

    Reviewed-by: Bhawanpreet Lakha <Bhawanpreet.Lakha@amd.com>
    Acked-by: Wayne Lin <wayne.lin@amd.com>
    Signed-off-by: Leo Li <sunpeng.li@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32c1dde16f0a5c8da320b9525bbf35b7c404da12
Author: Janis Schoetterl-Glausch <scgl@linux.ibm.com>
Date:   Thu Aug 25 21:25:40 2022 +0200

    KVM: s390: Pass initialized arg even if unused

    [ Upstream commit b3cefd6bf16e7234ffbd4209f6083060f4e35f59 ]

    This silences smatch warnings reported by kbuild bot:
    arch/s390/kvm/gaccess.c:859 guest_range_to_gpas() error: uninitialized symbol 'prot'.
    arch/s390/kvm/gaccess.c:1064 access_guest_with_key() error: uninitialized symbol 'prot'.

    This is because it cannot tell that the value is not used in this case.
    The trans_exc* only examine prot if code is PGM_PROTECTION.
    Pass a dummy value for other codes.

    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Janis Schoetterl-Glausch <scgl@linux.ibm.com>
    Reviewed-by: Claudio Imbrenda <imbrenda@linux.ibm.com>
    Link: https://lore.kernel.org/r/20220825192540.1560559-1-scgl@linux.ibm.com
    Signed-off-by: Christian Borntraeger <borntraeger@linux.ibm.com>
    Signed-off-by: Janosch Frank <frankja@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5b4e1c1ffba3d6a30990eace8e41b2ec8040c209
Author: Jianglei Nie <niejianglei2021@163.com>
Date:   Wed Sep 14 09:42:38 2022 +0800

    net: atlantic: fix potential memory leak in aq_ndev_close()

    [ Upstream commit 65e5d27df61283e5390f04b09dc79cd832f95607 ]

    If aq_nic_stop() fails, aq_ndev_close() returns err without calling
    aq_nic_deinit() to release the relevant memory and resource, which
    will lead to a memory leak.

    We can fix it by deleting the if condition judgment and goto statement to
    call aq_nic_deinit() directly after aq_nic_stop() to fix the memory leak.

    Signed-off-by: Jianglei Nie <niejianglei2021@163.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfaf7f090a4098a5780d805d70ac21f53b11775e
Author: David Gow <davidgow@google.com>
Date:   Wed Sep 21 14:48:55 2022 +0800

    arch: um: Mark the stack non-executable to fix a binutils warning

    [ Upstream commit bd71558d585ac61cfd799db7f25e78dca404dd7a ]

    Since binutils 2.39, ld will print a warning if any stack section is
    executable, which is the default for stack sections on files without a
    .note.GNU-stack section.

    This was fixed for x86 in commit ffcf9c5700e4 ("x86: link vdso and boot with -z
noexecstack --no-warn-rwx-segments"),
    but remained broken for UML, resulting in several warnings:

    /usr/bin/ld: warning: arch/x86/um/vdso/vdso.o: missing .note.GNU-stack section implies
executable stack
    /usr/bin/ld: NOTE: This behaviour is deprecated and will be removed in a future version
of the linker
    /usr/bin/ld: warning: .tmp_vmlinux.kallsyms1 has a LOAD segment with RWX permissions
    /usr/bin/ld: warning: .tmp_vmlinux.kallsyms1.o: missing .note.GNU-stack section implies
executable stack
    /usr/bin/ld: NOTE: This behaviour is deprecated and will be removed in a future version
of the linker
    /usr/bin/ld: warning: .tmp_vmlinux.kallsyms2 has a LOAD segment with RWX permissions
    /usr/bin/ld: warning: .tmp_vmlinux.kallsyms2.o: missing .note.GNU-stack section implies
executable stack
    /usr/bin/ld: NOTE: This behaviour is deprecated and will be removed in a future version
of the linker
    /usr/bin/ld: warning: vmlinux has a LOAD segment with RWX permissions

    Link both the VDSO and vmlinux with -z noexecstack, fixing the warnings
    about .note.GNU-stack sections. In addition, pass --no-warn-rwx-segments
    to dodge the remaining warnings about LOAD segments with RWX permissions
    in the kallsyms objects. (Note that this flag is apparently not
    available on lld, so hide it behind a test for BFD, which is what the
    x86 patch does.)

    Link: https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?
id=ffcf9c5700e49c0aee42dcba9a12ba21338e8136
    Link: https://sourceware.org/git/?p=binutils-
gdb.git;a=commit;h=ba951afb99912da01a6e8434126b8fac7aa75107
    Signed-off-by: David Gow <davidgow@google.com>
    Reviewed-by: Lukas Straub <lukasstraub2@web.de>
    Tested-by: Lukas Straub <lukasstraub2@web.de>
```

```
    Acked-by: Randy Dunlap <rdunlap@infradead.org> # build-tested
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a460f846d580952278793974849e2543b2aecbd5
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Thu Sep 15 22:32:54 2022 +0200

    gpio: ftgpio010: Make irqchip immutable

    [ Upstream commit ab637d48363d7b8ee67ae089808a8bc6051d53c4 ]

    This turns the FTGPIO010 irqchip immutable.

    Tested on the D-Link DIR-685.

    Cc: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Acked-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d71f03dbcc17e6ba4bc90daf6bdd7dad93a20764
Author: Lukas Straub <lukasstraub2@web.de>
Date:   Fri Aug 26 15:29:31 2022 +0000

    um: Cleanup compiler warning in arch/x86/um/tls_32.c

    [ Upstream commit d27fff3499671dc23a08efd01cdb8b3764a391c4 ]

    arch.tls_array is statically allocated so checking for NULL doesn't
    make sense. This causes the compiler warning below.

    Remove the checks to silence these warnings.

    ../arch/x86/um/tls_32.c: In function 'get_free_idx':
    ../arch/x86/um/tls_32.c:68:13: warning: the comparison will always evaluate as 'true'
for the address of 'tls_array' will never be NULL [-Waddress]
       68 |         if (!t->arch.tls_array)
          |             ^
    In file included from ../arch/x86/um/asm/processor.h:10,
                     from ../include/linux/rcupdate.h:30,
                     from ../include/linux/rculist.h:11,
                     from ../include/linux/pid.h:5,
                     from ../include/linux/sched.h:14,
                     from ../arch/x86/um/tls_32.c:7:
    ../arch/x86/um/asm/processor_32.h:22:31: note: 'tls_array' declared here
       22 |         struct uml_tls_struct tls_array[GDT_ENTRY_TLS_ENTRIES];
          |                               ^~~~~~~~~
    ../arch/x86/um/tls_32.c: In function 'get_tls_entry':
    ../arch/x86/um/tls_32.c:243:13: warning: the comparison will always evaluate as 'true'
for the address of 'tls_array' will never be NULL [-Waddress]
      243 |         if (!t->arch.tls_array)
          |             ^
    ../arch/x86/um/asm/processor_32.h:22:31: note: 'tls_array' declared here
       22 |         struct uml_tls_struct tls_array[GDT_ENTRY_TLS_ENTRIES];
          |                               ^~~~~~~~~

    Signed-off-by: Lukas Straub <lukasstraub2@web.de>
    Acked-by: Randy Dunlap <rdunlap@infradead.org> # build-tested
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7105d2a77a63eb36309ada16e3183c1fb45eda0
Author: Lukas Straub <lukasstraub2@web.de>
```

```
Date:   Fri Aug 26 15:29:27 2022 +0000

    um: Cleanup syscall_handler_t cast in syscalls_32.h

    [ Upstream commit 61670b4d270c71219def1fbc9441debc2ac2e6e9 ]

    Like in f4f03f299a56ce4d73c5431e0327b3b6cb55ebb9
    "um: Cleanup syscall_handler_t definition/cast, fix warning",
    remove the cast to to fix the compiler warning.

    Signed-off-by: Lukas Straub <lukasstraub2@web.de>
    Acked-by: Randy Dunlap <rdunlap@infradead.org> # build-tested
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9cdbc61abaa0611d9d2cb169586e7bb0daa128e1
Author: Jaroslav Kysela <perex@perex.cz>
Date:   Tue Sep 13 09:02:16 2022 +0200

    ALSA: hda/hdmi: Fix the converter reuse for the silent stream

    [ Upstream commit 5f80d6bd2b01de4cafac3302f58456bf860322fc ]

    When the user space pcm stream uses the silent stream converter,
    it is no longer allocated for the silent stream. Clear the appropriate
    flag in the hdmi_pcm_open() function. The silent stream setup may
    be applied in hdmi_pcm_close() (and the error path - open fcn) again.

    If the flag is not cleared, the reuse conditions for the silent
    stream converter in hdmi_choose_cvt() may improperly share
    this converter.

    Cc: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Signed-off-by: Jaroslav Kysela <perex@perex.cz>
    Link: https://lore.kernel.org/r/20220913070216.3233974-1-perex@perex.cz
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7480deff597a5b959ced9587622596c3c3773aa2
Author: Oleksandr Mazur <oleksandr.mazur@plvision.eu>
Date:   Thu Sep 8 16:14:46 2022 +0300

    net: marvell: prestera: add support for for Aldrin2

    [ Upstream commit 9124dbcc2dd6c51e81f97f63f7807118c4eb140a ]

    Aldrin2 (98DX8525) is a Marvell Prestera PP, with 100G support.

    Signed-off-by: Oleksandr Mazur <oleksandr.mazur@plvision.eu>

    V2:
      - retarget to net tree instead of net-next;
      - fix missed colon in patch subject ('net marvell' vs 'net: mavell');
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be248292a4836a5f110d422ccdbae9f04cb5574b
Author: Haimin Zhang <tcs.kernel@gmail.com>
Date:   Thu Sep 8 20:19:27 2022 +0800

    net/ieee802154: fix uninit value bug in dgram_sendmsg

    [ Upstream commit 94160108a70c8af17fa1484a37e05181c0e094af ]

    There is uninit value bug in dgram_sendmsg function in
```

```
    net/ieee802154/socket.c when the length of valid data pointed by the
    msg->msg_name isn't verified.

    We introducing a helper function ieee802154_sockaddr_check_size to
    check namelen. First we check there is addr_type in ieee802154_addr_sa.
    Then, we check namelen according to addr_type.

    Also fixed in raw_bind, dgram_bind, dgram_connect.

    Signed-off-by: Haimin Zhang <tcs_kernel@tencent.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce57f7b2e61404e28453cd75afa2a013c263c34f
Author: Letu Ren <fantasquex@gmail.com>
Date:   Fri Nov 12 20:06:41 2021 +0800

    scsi: qedf: Fix a UAF bug in __qedf_probe()

    [ Upstream commit fbfe96869b782364caebae0445763969ddb6ea67 ]

    In __qedf_probe(), if qedf->cdev is NULL which means
    qed_ops->common->probe() failed, then the program will goto label err1, and
    scsi_host_put() will free lport->host pointer. Because the memory qedf
    points to is allocated by libfc_host_alloc(), it will be freed by
    scsi_host_put(). However, the if statement below label err0 only checks
    whether qedf is NULL but doesn't check whether the memory has been freed.
    So a UAF bug can occur.

    There are two ways to reach the statements below err0. The first one is
    described as before, "qedf" should be set to NULL. The second one is goto
    "err0" directly. In the latter scenario qedf hasn't been changed and it has
    the initial value NULL. As a result the if statement is not reachable in
    any situation.

    The KASAN logs are as follows:

    [    2.312969] BUG: KASAN: use-after-free in __qedf_probe+0x5dcf/0x6bc0
    [    2.312969]
    [    2.312969] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.12.0-59-
gc9ba5276e321-prebuilt.qemu.org 04/01/2014
    [    2.312969] Call Trace:
    [    2.312969]  dump_stack_lvl+0x59/0x7b
    [    2.312969]  print_address_description+0x7c/0x3b0
    [    2.312969]  ? __qedf_probe+0x5dcf/0x6bc0
    [    2.312969]  __kasan_report+0x160/0x1c0
    [    2.312969]  ? __qedf_probe+0x5dcf/0x6bc0
    [    2.312969]  kasan_report+0x4b/0x70
    [    2.312969]  ? kobject_put+0x25d/0x290
    [    2.312969]  kasan_check_range+0x2ca/0x310
    [    2.312969]  __qedf_probe+0x5dcf/0x6bc0
    [    2.312969]  ? selinux_kernfs_init_security+0xdc/0x5f0
    [    2.312969]  ? trace_rpm_return_int_rcuidle+0x18/0x120
    [    2.312969]  ? rpm_resume+0xa5c/0x16e0
    [    2.312969]  ? qedf_get_generic_tlv_data+0x160/0x160
    [    2.312969]  local_pci_probe+0x13c/0x1f0
    [    2.312969]  pci_device_probe+0x37e/0x6c0

    Link: https://lore.kernel.org/r/20211112120641.16073-1-fantasquex@gmail.com
    Reported-by: Zheyu Ma <zheyuma97@gmail.com>
    Acked-by: Saurav Kashyap <skashyap@marvell.com>
    Co-developed-by: Wende Tan <twd2.me@gmail.com>
    Signed-off-by: Wende Tan <twd2.me@gmail.com>
    Signed-off-by: Letu Ren <fantasquex@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
```

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

commit 5f2f005094963accab7c01f66b439f969a050a45
Author: Yifan Zhang <yifan1.zhang@amd.com>
Date:   Tue Aug 30 23:18:47 2022 +0800

    drm/amdgpu/mes: zero the sdma_hqd_mask of 2nd SDMA engine for SDMA 6.0.1

    [ Upstream commit 0af4ed0c329ebb4cef95fda4fcdbfcdea0255442 ]

    there is only one SDMA engine in SDMA 6.0.1, the sdma_hqd_mask has to be
    zeroed for the 2nd engine, otherwise MES scheduler will consider 2nd
    engine exists and map/unmap SDMA queues to the non-existent engine.

    Signed-off-by: Yifan Zhang <yifan1.zhang@amd.com>
    Reviewed-by: Tim Huang <Tim.Huang@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8895cfb48b004c30470e757f2022d3416dd199d
Author: Sergei Antonov <saproj@gmail.com>
Date:   Wed Sep 7 20:53:41 2022 +0300

    ARM: dts: fix Moxa SDIO 'compatible', remove 'sdhci' misnomer

    [ Upstream commit 02181e68275d28cab3c3f755852770367f1bc229 ]

    Driver moxart-mmc.c has .compatible = "moxa,moxart-mmc".

    But moxart .dts/.dtsi and the documentation file moxa,moxart-dma.txt
    contain compatible = "moxa,moxart-sdhci".

    Change moxart .dts/.dtsi files and moxa,moxart-dma.txt to match the driver.

    Replace 'sdhci' with 'mmc' in names too, since SDHCI is a different
    controller from FTSDC010.

    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sergei Antonov <saproj@gmail.com>
    Cc: Jonas Jensen <jonas.jensen@gmail.com>
    Link: https://lore.kernel.org/r/20220907175341.1477383-1-saproj@gmail.com'
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e834ee927a512fe0901f46f58622bef0d6e84bd
Author: Jason A. Donenfeld <Jason@zx2c4.com>
Date:   Mon Sep 5 19:22:46 2022 +0200

    wifi: iwlwifi: don't spam logs with NSS>2 messages

    [ Upstream commit 4d8421f2dd88583cc7a4d6c2a5532c35e816a52a ]

    I get a log line like this every 4 seconds when connected to my AP:

    [15650.221468] iwlwifi 0000:09:00.0: Got NSS = 4 - trimming to 2

    Looking at the code, this seems to be related to a hardware limitation,
    and there's nothing to be done. In an effort to keep my dmesg
    manageable, downgrade this error to "debug" rather than "info".

    Cc: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Jason A. Donenfeld <Jason@zx2c4.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20220905172246.105383-1-Jason@zx2c4.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 229769759ef817596cf295b38fa2e04e7a62c719
Author: Swati Agarwal <swati.agarwal@xilinx.com>
Date:   Wed Aug 17 11:41:25 2022 +0530

    dmaengine: xilinx_dma: Report error in case of dma_set_mask_and_coherent API failure

    [ Upstream commit 8f2b6bc79c32f0fa60df000ae387a790ec80eae9 ]

    The driver does not handle the failure case while calling
    dma_set_mask_and_coherent API.

    In case of failure, capture the return value of API and then report an
    error.

    Addresses-coverity: Unchecked return value (CHECKED_RETURN)

    Signed-off-by: Swati Agarwal <swati.agarwal@xilinx.com>
    Reviewed-by: Radhey Shyam Pandey <radhey.shyam.pandey@xilinx.com>
    Link: https://lore.kernel.org/r/20220817061125.4720-4-swati.agarwal@xilinx.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e38f675825679c7e09bf31fcbddcbe9d95a7b294
Author: Swati Agarwal <swati.agarwal@xilinx.com>
Date:   Wed Aug 17 11:41:24 2022 +0530

    dmaengine: xilinx_dma: cleanup for fetching xlnx,num-fstores property

    [ Upstream commit 462bce790e6a7e68620a4ce260cc38f7ed0255d5 ]

    Free the allocated resources for missing xlnx,num-fstores property.

    Signed-off-by: Swati Agarwal <swati.agarwal@xilinx.com>
    Link: https://lore.kernel.org/r/20220817061125.4720-3-swati.agarwal@xilinx.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92f8963831f164573f7fb1265ebee9478c99cf68
Author: Swati Agarwal <swati.agarwal@xilinx.com>
Date:   Wed Aug 17 11:41:23 2022 +0530

    dmaengine: xilinx_dma: Fix devm_platform_ioremap_resource error handling

    [ Upstream commit 91df7751eb890e970afc08f50b8f0fa5ea39e03d ]

    Add missing cleanup in devm_platform_ioremap_resource().
    When probe fails remove dma channel resources and disable clocks in
    accordance with the order of resources allocated .

    Signed-off-by: Swati Agarwal <swati.agarwal@xilinx.com>
    Link: https://lore.kernel.org/r/20220817061125.4720-2-swati.agarwal@xilinx.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 666f5be3593c5795a1bbaeaa6cf0037ae1a23eb2
Author: Frank Wunderlich <frank-w@public-files.de>
Date:   Sun Aug 21 14:19:29 2022 +0200

    arm64: dts: rockchip: fix upper usb port on BPI-R2-Pro

    [ Upstream commit 388f9f0a7ff84b7890a24499a3a1fea0cad21373 ]

    - extcon is no more needed in 5.19 - so drop it
      commit 51a9b2c03dd3 ("phy: rockchip-inno-usb2: Handle ID IRQ")
```

```
      - dr_mode was changed from host to otg in rk356x.dtsi
        commit bc405bb3eeee ("arm64: dts: rockchip: enable otg/drd
          operation of usb_host0_xhci in rk356x")
        change it back on board level as id-pin on r2pro is not connected

      Signed-off-by: Frank Wunderlich <frank-w@public-files.de>
      Link: https://lore.kernel.org/r/20220821121929.244112-1-linux@fw-web.de
      Signed-off-by: Heiko Stuebner <heiko@sntech.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit edc5c66d15f96617e2deeb194eb73f13732a3443
  Author: Cristian Marussi <cristian.marussi@arm.com>
  Date:   Wed Aug 17 18:27:31 2022 +0100

      firmware: arm_scmi: Add SCMI PM driver remove routine

      [ Upstream commit dea796fcab0a219830831c070b8dc367d7e0f708 ]

      Currently, when removing the SCMI PM driver not all the resources
      registered with genpd subsystem are properly de-registered.

      As a side effect of this after a driver unload/load cycle you get a
      splat with a few warnings like this:

       | debugfs: Directory 'BIG_CPU0' with parent 'pm_genpd' already present!
       | debugfs: Directory 'BIG_CPU1' with parent 'pm_genpd' already present!
       | debugfs: Directory 'LITTLE_CPU0' with parent 'pm_genpd' already present!
       | debugfs: Directory 'LITTLE_CPU1' with parent 'pm_genpd' already present!
       | debugfs: Directory 'LITTLE_CPU2' with parent 'pm_genpd' already present!
       | debugfs: Directory 'LITTLE_CPU3' with parent 'pm_genpd' already present!
       | debugfs: Directory 'BIG_SSTOP' with parent 'pm_genpd' already present!
       | debugfs: Directory 'LITTLE_SSTOP' with parent 'pm_genpd' already present!
       | debugfs: Directory 'DBGSYS' with parent 'pm_genpd' already present!
       | debugfs: Directory 'GPUTOP' with parent 'pm_genpd' already present!

      Add a proper scmi_pm_domain_remove callback to the driver in order to
      take care of all the needed cleanups not handled by devres framework.

      Link: https://lore.kernel.org/r/20220817172731.1185305-7-cristian.marussi@arm.com
      Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
      Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 8e880f30946fd4d8a864defeeada8a9a5b48848a
  Author: Cristian Marussi <cristian.marussi@arm.com>
  Date:   Wed Aug 17 18:27:28 2022 +0100

      firmware: arm_scmi: Harden accesses to the sensor domains

      [ Upstream commit 76f89c954788763db575fb512a40bd483864f1e9 ]

      Accessing sensor domains descriptors by the index upon the SCMI drivers
      requests through the SCMI sensor operations interface can potentially
      lead to out-of-bound violations if the SCMI driver misbehave.

      Add an internal consistency check before any such domains descriptors
      accesses.

      Link: https://lore.kernel.org/r/20220817172731.1185305-4-cristian.marussi@arm.com
      Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
      Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 41bf1b0ad95395ec6d5172fb58540323788d9397
  Author: Cristian Marussi <cristian.marussi@arm.com>
```

```
Date:   Wed Aug 17 18:27:27 2022 +0100

    firmware: arm_scmi: Improve checks in the info_get operations

    [ Upstream commit 1ecb7d27b1af6705e9a4e94415b4d8cc8cf2fbfb ]

    SCMI protocols abstract and expose a number of protocol specific
    resources like clocks, sensors and so on. Information about such
    specific domain resources are generally exposed via an `info_get`
    protocol operation.

    Improve the sanity check on these operations where needed.

    Link: https://lore.kernel.org/r/20220817172731.1185305-3-cristian.marussi@arm.com
    Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 81de80330fa6907aec32eb54c5619059e6e36452
Author: Dongliang Mu <mudongliangabcd@gmail.com>
Date:   Tue Aug 16 12:08:58 2022 +0800

    fs: fix UAF/GPF bug in nilfs_mdt_destroy

    commit 2e488f13755ffbb60f307e991b27024716a33b29 upstream.

    In alloc_inode, inode_init_always() could return -ENOMEM if
    security_inode_alloc() fails, which causes inode->i_private
    uninitialized. Then nilfs_is_metadata_file_inode() returns
    true and nilfs_free_inode() wrongly calls nilfs_mdt_destroy(),
    which frees the uninitialized inode->i_private
    and leads to crashes(e.g., UAF/GPF).

    Fix this by moving security_inode_alloc just prior to
    this_cpu_inc(nr_inodes)


Link: https://lkml.kernel.org/r/CAFcO6XOcf1Jj2SeGt=jJV59wmhESeSKpfR0omdFRq+J9nD1vfQ@mail.gma
il.com
    Reported-by: butt3rflyh4ck <butterflyhuangxx@gmail.com>
    Reported-by: Hao Sun <sunhao.th@gmail.com>
    Reported-by: Jiacheng Xu <stitch@zju.edu.cn>
    Reviewed-by: Christian Brauner (Microsoft) <brauner@kernel.org>
    Signed-off-by: Dongliang Mu <mudongliangabcd@gmail.com>
    Cc: Al Viro <viro@zeniv.linux.org.uk>
    Cc: stable@vger.kernel.org
    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8529dfb5ef724a2edb1c1b3567688376b3db441d
Author: Jalal Mostafa <jalal.a.mostapha@gmail.com>
Date:   Wed Sep 21 13:57:01 2022 +0000

    xsk: Inherit need_wakeup flag for shared sockets

    commit 60240bc26114543fcbfcd8a28466e67e77b20388 upstream.

    The flag for need_wakeup is not set for xsks with `XDP_SHARED_UMEM`
    flag and of different queue ids and/or devices. They should inherit
    the flag from the first socket buffer pool since no flags can be
    specified once `XDP_SHARED_UMEM` is specified.

    Fixes: b5aea28dca134 ("xsk: Add shared umem support between queue ids")
    Signed-off-by: Jalal Mostafa <jalal.a.mostapha@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
```

```
    Acked-by: Magnus Karlsson <magnus.karlsson@intel.com>
    Link: https://lore.kernel.org/bpf/20220921135701.10199-1-jalal.a.mostapha@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7ebe331054cfb8ad97c8542e46d10a3887be1e01
Author: Shuah Khan <skhan@linuxfoundation.org>
Date:   Thu Sep 1 15:23:19 2022 -0600

    docs: update mediator information in CoC docs

    commit 8bfdfa0d6b929ede7b6189e0e546ceb6a124d05d upstream.

    Update mediator information in the CoC interpretation document.

    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Link: https://lore.kernel.org/r/20220901212319.56644-1-skhan@linuxfoundation.org
    Cc: stable@vger.kernel.org
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b760c94de19c8296a6aa861845b348f14f335f7
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Sep 29 22:57:43 2022 -0700

    hardening: Remove Clang's enable flag for -ftrivial-auto-var-init=zero

    commit 607e57c6c62c00965ae276902c166834ce73014a upstream.

    Now that Clang's -enable-trivial-auto-var-init-zero-knowing-it-will-be-removed-from-
clang
    option is no longer required, remove it from the command line. Clang 16
    and later will warn when it is used, which will cause Kconfig to think
    it can't use -ftrivial-auto-var-init=zero at all. Check for whether it
    is required and only use it when so.

    Cc: Nathan Chancellor <nathan@kernel.org>
    Cc: Masahiro Yamada <masahiroy@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: linux-kbuild@vger.kernel.org
    Cc: llvm@lists.linux.dev
    Cc: stable@vger.kernel.org
    Fixes: f02003c860d9 ("hardening: Avoid harmless Clang option under
CONFIG_INIT_STACK_ALL_ZERO")
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 19a4cb1c4eb8f15a72700a20f778c52028e687b7
Author: Sami Tolvanen <samitolvanen@google.com>
Date:   Fri Sep 30 20:33:10 2022 +0000

    Makefile.extrawarn: Move -Wcast-function-type-strict to W=1

    commit 2120635108b35ecad9c59c8b44f6cbdf4f98214e upstream.

    We enable -Wcast-function-type globally in the kernel to warn about
    mismatching types in function pointer casts. Compilers currently
    warn only about ABI incompability with this flag, but Clang 16 will
    enable a stricter version of the check by default that checks for an
    exact type match. This will be very noisy in the kernel, so disable
    -Wcast-function-type-strict without W=1 until the new warnings have
    been addressed.

    Cc: stable@vger.kernel.org
    Link: https://reviews.llvm.org/D134831
    Link: https://github.com/ClangBuiltLinux/linux/issues/1724
```

```
    Suggested-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Sami Tolvanen <samitolvanen@google.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20220930203310.4010564-1-samitolvanen@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8dc610c9478ccb0094d70663b345e46c16dbd5e7
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Tue Aug 30 13:58:42 2022 -0700

    sparc: Unbreak the build

    commit 17006e86a7641fa3c50324cfb602f0e74dac8527 upstream.

    Fix the following build errors:

    arch/sparc/mm/srmmu.c: In function ‘smp_flush_page_for_dma’:
    arch/sparc/mm/srmmu.c:1639:13: error: cast between incompatible function types from
‘void (*)(long unsigned int)’ to ‘void (*)(long unsigned int,  long unsigned int,  long
unsigned int,  long unsigned int,  long unsigned int)’ [-Werror=cast-function-type]
     1639 |         xc1((smpfunc_t) local_ops->page_for_dma, page);
          |             ^
    arch/sparc/mm/srmmu.c: In function ‘smp_flush_cache_mm’:
    arch/sparc/mm/srmmu.c:1662:29: error: cast between incompatible function types from
‘void (*)(struct mm_struct *)’ to ‘void (*)(long unsigned int,  long unsigned int,  long
unsigned int,  long unsigned int,  long unsigned int)’ [-Werror=cast-function-type]
     1662 |                         xc1((smpfunc_t) local_ops->cache_mm, (unsigned long)
mm);
          |
    [ ... ]

    Compile-tested only.

    Fixes: 552a23a0e5d0 ("Makefile: Enable -Wcast-function-type")
    Cc: stable@vger.kernel.org
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Tested-by: Andreas Larsson <andreas@gaisler.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20220830205854.1918026-1-bvanassche@acm.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```