```
commit f86aa267e180b23d8d24ee6886bd3bf9a360112e
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Mon Jul 19 10:04:55 2021 +0200

    Linux 5.13.3

    Link: https://lore.kernel.org/r/20210715182613.933608881@linuxfoundation.org
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Link: https://lore.kernel.org/r/20210716182150.239646976@linuxfoundation.org
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e472b276a0d2180808009be38105e12754432e2a
Author: Chao Yu <chao@kernel.org>
Date:   Fri May 7 18:10:38 2021 +0800

    f2fs: fix to avoid racing on fsync_entry_slab by multi filesystem instances

    commit cad83c968c2ebe97905f900326988ed37146c347 upstream.

    As syzbot reported, there is an use-after-free issue during f2fs recovery:

    Use-after-free write at 0xffff88823bc16040 (in kfence-#10):
     kmem_cache_destroy+0x1f/0x120 mm/slab_common.c:486
     f2fs_recover_fsync_data+0x75b0/0x8380 fs/f2fs/recovery.c:869
     f2fs_fill_super+0x9393/0xa420 fs/f2fs/super.c:3945
     mount_bdev+0x26c/0x3a0 fs/super.c:1367
     legacy_get_tree+0xea/0x180 fs/fs_context.c:592
     vfs_get_tree+0x86/0x270 fs/super.c:1497
     do_new_mount fs/namespace.c:2905 [inline]
     path_mount+0x196f/0x2be0 fs/namespace.c:3235
     do_mount fs/namespace.c:3248 [inline]
     __do_sys_mount fs/namespace.c:3456 [inline]
     __se_sys_mount+0x2f9/0x3b0 fs/namespace.c:3433
     do_syscall_64+0x3f/0xb0 arch/x86/entry/common.c:47
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    The root cause is multi f2fs filesystem instances can race on accessing
    global fsync_entry_slab pointer, result in use-after-free issue of slab
    cache, fixes to init/destroy this slab cache only once during module
    init/destroy procedure to avoid this issue.

    Reported-by: syzbot+9d90dad32dd9727ed084@syzkaller.appspotmail.com
    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 258fd821f69378453c071b9dd767b298810fc766
Author: Tetsuo Handa <penguin-kernel@i-love.sakura.ne.jp>
Date:   Mon Apr 12 22:45:50 2021 +0900

    smackfs: restrict bytes count in smk_set_cipso()

    commit 49ec114a6e62d8d320037ce71c1aaf9650b3cafd upstream.

    Oops, I failed to update subject line.

    From 07571157c91b98ce1a4aa70967531e64b78e8346 Mon Sep 17 00:00:00 2001
```

```
    From: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Date: Mon, 12 Apr 2021 22:25:06 +0900
    Subject: smackfs: restrict bytes count in smk_set_cipso()

    Commit 7ef4c19d245f3dc2 ("smackfs: restrict bytes count in smackfs write
    functions") missed that count > SMK_CIPSOMAX check applies to only
    format == SMK_FIXED24_FMT case.

    Reported-by: syzbot <syzbot+77c53db50c9fff774e8e@syzkaller.appspotmail.com>
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Signed-off-by: Casey Schaufler <casey@schaufler-ca.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bfb48b54db25c3b4ef4bef5e0691464ebc4aa335
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Jun 14 12:34:03 2021 +0200

    media: v4l2-core: explicitly clear ioctl input data

    commit 7b53cca764f9b291b7907fcd39d9e66ad728ee0b upstream.

    As seen from a recent syzbot bug report, mistakes in the compat ioctl
    implementation can lead to uninitialized kernel stack data getting used
    as input for driver ioctl handlers.

    The reported bug is now fixed, but it's possible that other related
    bugs are still present or get added in the future. As the drivers need
    to check user input already, the possible impact is fairly low, but it
    might still cause an information leak.

    To be on the safe side, always clear the entire ioctl buffer before
    calling the conversion handler functions that are meant to initialize
    them.

    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 42f102ea1943ecb10a0756bf75424de5d1d5beed
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Sun Jun 6 17:24:05 2021 +0300

    jfs: fix GPF in diFree

    commit 9d574f985fe33efd6911f4d752de6f485a1ea732 upstream.

    Avoid passing inode with
    JFS_SBI(inode->i_sb)->ipimap == NULL to
    diFree()[1]. GFP will appear:

            struct inode *ipimap = JFS_SBI(ip->i_sb)->ipimap;
            struct inomap *imap = JFS_IP(ipimap)->i_imap;

    JFS_IP() will return invalid pointer when ipimap == NULL

    Call Trace:
     diFree+0x13d/0x2dc0 fs/jfs/jfs_imap.c:853 [1]
     jfs_evict_inode+0x2c9/0x370 fs/jfs/inode.c:154
     evict+0x2ed/0x750 fs/inode.c:578
     iput_final fs/inode.c:1654 [inline]
     iput.part.0+0x3fe/0x820 fs/inode.c:1680
     iput+0x58/0x70 fs/inode.c:1670
```

```
    Reported-and-tested-by: syzbot+0a89a7b56db04c21a656@syzkaller.appspotmail.com
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Dave Kleikamp <dave.kleikamp@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7ed572cdf11081f8f9e07abd4bea56a3f2c4edbd
Author: Theodore Ts'o <tytso@mit.edu>
Date:   Fri Jul 2 12:45:02 2021 -0400

    ext4: fix possible UAF when remounting r/o a mmp-protected file system

    commit 61bb4a1c417e5b95d9edb4f887f131de32e419cb upstream.

    After commit 618f003199c6 ("ext4: fix memory leak in
    ext4_fill_super"), after the file system is remounted read-only, there
    is a race where the kmmpd thread can exit, causing sbi->s_mmp_tsk to
    point at freed memory, which the call to ext4_stop_mmpd() can trip
    over.

    Fix this by only allowing kmmpd() to exit when it is stopped via
    ext4_stop_mmpd().

    Link: https://lore.kernel.org/r/20210707002433.3719773-1-tytso@mit.edu
    Reported-by: Ye Bin <yebin10@huawei.com>
    Bug-Report-Link: <20210629143603.2166962-1-yebin10@huawei.com>
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Reviewed-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e9f6b57ef11459329367edd7c8c58a555886c4c8
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Thu Apr 29 12:50:57 2021 +0200

    drm/ast: Remove reference to struct drm_device.pdev

    commit 0ecb51824e838372e01330752503ddf9c0430ef7 upstream.

    Using struct drm_device.pdev is deprecated. Upcast with to_pci_dev()
    from struct drm_device.dev to get the PCI device structure.

    v9:
            * fix remaining pdev references

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Reviewed-by: Michael J. Ruhl <michael.j.ruhl@intel.com>
    Fixes: ba4e0339a6a3 ("drm/ast: Fixed CVE for DP501")
    Cc: KuoHsiang Chou <kuohsiang_chou@aspeedtech.com>
    Cc: kernel test robot <lkp@intel.com>
    Cc: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: Dave Airlie <airlied@redhat.com>
    Cc: dri-devel@lists.freedesktop.org
    Link: https://patchwork.freedesktop.org/patch/msgid/20210429105101.25667-2-
tzimmermann@suse.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7eefc573bcbf4a1b035e43b571923ef570f0361f
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue Jun 8 14:34:08 2021 +0800

    pinctrl: mcp23s08: Fix missing unlock on error in mcp23s08_irq()

    commit 884af72c90016cfccd5717439c86b48702cbf184 upstream.

    Add the missing unlock before return from function mcp23s08_irq()
    in the error handling case.
```

```
    v1-->v2:
       remove the "return IRQ_HANDLED" line

    Fixes: 897120d41e7a ("pinctrl: mcp23s08: fix race condition in irq handler")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Link: https://lore.kernel.org/r/1623134048-56051-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e294f97e8af8542cf4e9c58f09de463b7c03e641
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Mon Jun 21 14:48:29 2021 -0400

    dm writecache: write at least 4k when committing

    commit 867de40c4c23e6d7f89f9ce4272a5d1b1484c122 upstream.

    SSDs perform badly with sub-4k writes (because they perfrorm
    read-modify-write internally), so make sure writecache writes at least
    4k when committing.

    Fixes: 991bd8d7bc78 ("dm writecache: commit just one block, not a full page")
    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 691e71cf1fd19e645ec632aee6ae5183112dacfa
Author: Benjamin Drung <bdrung@posteo.de>
Date:   Sat Jun 5 22:15:36 2021 +0200

    media: uvcvideo: Fix pixel format change for Elgato Cam Link 4K

    commit 4c6e0976295add7f0ed94d276c04a3d6f1ea8f83 upstream.

    The Elgato Cam Link 4K HDMI video capture card reports to support three
    different pixel formats, where the first format depends on the connected
    HDMI device.

    ```
    $ v4l2-ctl -d /dev/video0 --list-formats-ext
    ioctl: VIDIOC_ENUM_FMT
            Type: Video Capture

            [0]: 'NV12' (Y/CbCr 4:2:0)
                    Size: Discrete 3840x2160
                            Interval: Discrete 0.033s (29.970 fps)
            [1]: 'NV12' (Y/CbCr 4:2:0)
                    Size: Discrete 3840x2160
                            Interval: Discrete 0.033s (29.970 fps)
            [2]: 'YU12' (Planar YUV 4:2:0)
                    Size: Discrete 3840x2160
                            Interval: Discrete 0.033s (29.970 fps)
    ```

    Changing the pixel format to anything besides the first pixel format
    does not work:

    ```
    $ v4l2-ctl -d /dev/video0 --try-fmt-video pixelformat=YU12
    Format Video Capture:
            Width/Height      : 3840/2160
            Pixel Format      : 'NV12' (Y/CbCr 4:2:0)
            Field             : None
```
```

```
            Bytes per Line    : 3840
            Size Image        : 12441600
            Colorspace        : sRGB
            Transfer Function : Rec. 709
            YCbCr/HSV Encoding: Rec. 709
            Quantization      : Default (maps to Limited Range)
            Flags             :
```

    User space applications like VLC might show an error message on the
    terminal in that case:

    ```
    libv4l2: error set_fmt gave us a different result than try_fmt!
    ```

    Depending on the error handling of the user space applications, they
    might display a distorted video, because they use the wrong pixel format
    for decoding the stream.

    The Elgato Cam Link 4K responds to the USB video probe
    VS_PROBE_CONTROL/VS_COMMIT_CONTROL with a malformed data structure: The
    second byte contains bFormatIndex (instead of being the second byte of
    bmHint). The first byte is always zero. The third byte is always 1.

    The firmware bug was reported to Elgato on 2020-12-01 and it was
    forwarded by the support team to the developers as feature request.
    There is no firmware update available since then. The latest firmware
    for Elgato Cam Link 4K as of 2021-03-23 has MCU 20.02.19 and FPGA 67.

    Therefore correct the malformed data structure for this device. The
    change was successfully tested with VLC, OBS, and Chromium using
    different pixel formats (YUYV, NV12, YU12), resolutions (3840x2160,
    1920x1080), and frame rates (29.970 and 59.940 fps).

    Cc: stable@vger.kernel.org
    Signed-off-by: Benjamin Drung <bdrung@posteo.de>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7c148146e07571acc6663aa42e66e24042447936
Author: Johan Hovold <johan@kernel.org>
Date: Mon May 24 13:09:19 2021 +0200

    media: gspca/sunplus: fix zero-length control requests

    commit b4bb4d425b7b02424afea2dfdcd77b3b4794175e upstream.

    The direction of the pipe argument must match the request-type direction
    bit or control requests may fail depending on the host-controller-driver
    implementation.

    Control transfers without a data stage are treated as OUT requests by
    the USB stack and should be using usb_sndctrlpipe(). Failing to do so
    will now trigger a warning.

    Fix the single zero-length control request which was using the
    read-register helper, and update the helper so that zero-length reads
    fail with an error message instead.

    Fixes: 6a7eba24e4f0 ("V4L/DVB (8157): gspca: all subdrivers")
    Cc: stable@vger.kernel.org   # 2.6.27
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>

```
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8f6a609903a04769d352a2e1362d8c85a1c59003
Author: Johan Hovold <johan@kernel.org>
Date:   Fri May 21 15:28:39 2021 +0200

    media: gspca/sq905: fix control-request direction

    commit 53ae298fde7adcc4b1432bce2dbdf8dac54dfa72 upstream.

    The direction of the pipe argument must match the request-type direction
    bit or control requests may fail depending on the host-controller-driver
    implementation.

    Fix the USB_REQ_SYNCH_FRAME request which erroneously used
    usb_sndctrlpipe().

    Fixes: 27d35fc3fb06 ("V4L/DVB (10639): gspca - sq905: New subdriver.")
    Cc: stable@vger.kernel.org      # 2.6.30
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 90be28e26596327684e19bb35e630646a7d84406
Author: Bernhard Wimmer <be.wimm@gmail.com>
Date:   Wed Apr 21 23:33:20 2021 +0200

    media: ccs: Fix the op_pll_multiplier address

    commit 0e3e0c9369c822b7f1dd11504eeb98cfd4aabf24 upstream.

    According to the CCS spec the op_pll_multiplier address is 0x030e,
    not 0x031e.

    Signed-off-by: Bernhard Wimmer <be.wimm@gmail.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Cc: stable@vger.kernel.org
    Fixes: 6493c4b777c2 ("media: smiapp: Import CCS definitions")
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d69b39d89f362cfeeb54a68690768d0d257b2c8f
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Mon May 17 21:18:14 2021 +0200

    media: zr364xx: fix memory leak in zr364xx_start_readpipe

    commit 0a045eac8d0427b64577a24d74bb8347c905ac65 upstream.

    syzbot reported memory leak in zr364xx driver.
    The problem was in non-freed urb in case of
    usb_submit_urb() fail.

    backtrace:
      [<ffffffff82baedf6>] kmalloc include/linux/slab.h:561 [inline]
      [<ffffffff82baedf6>] usb_alloc_urb+0x66/0xe0 drivers/usb/core/urb.c:74
      [<ffffffff82f7cce8>] zr364xx_start_readpipe+0x78/0x130
drivers/media/usb/zr364xx/zr364xx.c:1022
      [<ffffffff84251dfc>] zr364xx_board_init drivers/media/usb/zr364xx/zr364xx.c:1383
[inline]
      [<ffffffff84251dfc>] zr364xx_probe+0x6a3/0x851
drivers/media/usb/zr364xx/zr364xx.c:1516
      [<ffffffff82bb6507>] usb_probe_interface+0x177/0x370 drivers/usb/core/driver.c:396
```

```
       [<ffffffff826018a9>] really_probe+0x159/0x500 drivers/base/dd.c:576

    Fixes: ccbf035ae5de ("V4L/DVB (12278): zr364xx: implement V4L2_CAP_STREAMING")
    Cc: stable@vger.kernel.org
    Reported-by: syzbot+af4fa391ef18efdd5f69@syzkaller.appspotmail.com
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2f920c1543f7c253a6bfc21d679561be51a959c6
Author: Johan Hovold <johan@kernel.org>
Date:   Fri May 21 15:28:38 2021 +0200

    media: dtv5100: fix control-request directions

    commit 8c8b9a9be2afa8bd6a72ad1130532baab9fab89d upstream.

    The direction of the pipe argument must match the request-type direction
    bit or control requests may fail depending on the host-controller-driver
    implementation.

    Fix the control requests which erroneously used usb_rcvctrlpipe().

    Fixes: 8466028be792 ("V4L/DVB (8734): Initial support for AME DTV-5100 USB2.0 DVB-T")
    Cc: stable@vger.kernel.org      # 2.6.28
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7be5e037e564bd2479d414fe1198dc3e8652d278
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Fri Apr 23 17:19:11 2021 +0200

    media: i2c: ccs-core: fix pm_runtime_get_sync() usage count

    commit da3a1858c3a37c09446e1470c48352897d59d11b upstream.

    The pm_runtime_get_sync() internally increments the
    dev->power.usage_count without decrementing it, even on errors.

    There is a bug at ccs_pm_get_init(): when this function returns
    an error, the stream is not started, and RPM usage_count
    should not be incremented. However, if the calls to
    v4l2_ctrl_handler_setup() return errors, it will be kept
    incremented.

    At ccs_suspend() the best is to replace it by the new
    pm_runtime_resume_and_get(), introduced by:
    commit dd8088d5a896 ("PM: runtime: Add pm_runtime_resume_and_get to deal with usage
counter")
    in order to properly decrement the usage counter automatically,
    in the case of errors.

    Fixes: 96e3a6b92f23 ("media: smiapp: Avoid maintaining power state information")
    Cc: stable@vger.kernel.org
    Acked-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8b39b430dc0a16062d2c9a40b7be0da9bfd64edf
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Jun 14 12:34:09 2021 +0200
```

```
    media: subdev: disallow ioctl for saa6588/davinci

    commit 0a7790be182d32b9b332a37cb4206e24fe94b728 upstream.

    The saa6588_ioctl() function expects to get called from other kernel
    functions with a 'saa6588_command' pointer, but I found nothing stops it
    from getting called from user space instead, which seems rather dangerous.

    The same thing happens in the davinci vpbe driver with its VENC_GET_FLD
    command.

    As a quick fix, add a separate .command() callback pointer for this
    driver and change the two callers over to that.  This change can easily
    get backported to stable kernels if necessary, but since there are only
    two drivers, we may want to eventually replace this with a set of more
    specialized callbacks in the long run.

    Fixes: c3fda7f835b0 ("V4L/DVB (10537): saa6588: convert to v4l2_subdev.")
    Cc: stable@vger.kernel.org
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eda818ae5d3916a4685998555222fcc0af8020fb
Author: Pali Rohár <pali@kernel.org>
Date:   Fri Jun 25 00:26:20 2021 +0200

    PCI: aardvark: Implement workaround for the readback value of VEND_ID

    commit 7f71a409fe3d9358da07c77f15bb5b7960f12253 upstream.

    Marvell Armada 3700 Functional Errata, Guidelines, and Restrictions
    document describes in erratum 4.1 PCIe value of vendor ID (Ref #: 243):

        The readback value of VEND_ID (RD0070000h [15:0]) is 1B4Bh, while it
        should read 11ABh.

        The firmware can write the correct value, 11ABh, through VEND_ID
        (RD0076044h [15:0]).

    Implement this workaround in aardvark driver for both PCI vendor id and PCI
    subsystem vendor id.

    This change affects and fixes PCI vendor id of emulated PCIe root bridge.
    After this change emulated PCIe root bridge has correct vendor id.

    Link: https://lore.kernel.org/r/20210624222621.4776-5-pali@kernel.org
    Fixes: 8a3ebd8de328 ("PCI: aardvark: Implement emulated root PCI bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Marek Behún <kabel@kernel.org>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 678d93d3d2434fbf585f28f34693957e60519ac1
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Jun 24 23:33:43 2021 +0200

    PCI: aardvark: Fix checking for PIO Non-posted Request

    commit 8ceeac307a79f68c0d0c72d6e48b82fa424204ec upstream.

    PIO_NON_POSTED_REQ for PIO_STAT register is incorrectly defined. Bit 10 in
```

```
    register PIO_STAT indicates the response is to a non-posted request.

    Link: https://lore.kernel.org/r/20210624213345.3617-2-pali@kernel.org
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Marek Behún <kabel@kernel.org>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9789b16dc4e337fc8d1fffee25769b44eda6da87
Author: Konstantin Kharlamov <Hi-Angel@yandex.ru>
Date:   Fri May 21 02:55:01 2021 +0300

    PCI: Leave Apple Thunderbolt controllers on for s2idle or standby

    commit 4694ae373dc2114f9a82f6ae15737e65af0c6dea upstream.

    On Macbook 2013, resuming from suspend-to-idle or standby resulted in the
    external monitor no longer being detected, a stacktrace, and errors like
    this in dmesg:

      pcieport 0000:06:00.0: can't change power state from D3hot to D0 (config space
inaccessible)

    The reason is that we know how to turn power to the Thunderbolt controller
    *off* via the SXIO/SXFP/SXLF methods, but we don't know how to turn power
    back on.  We have to rely on firmware to turn the power back on.

    When going to the "suspend-to-idle" or "standby" system sleep states,
    firmware is not involved either on the suspend side or the resume side, so
    we can't use SXIO/SXFP/SXLF to turn the power off.

    Skip SXIO/SXFP/SXLF when firmware isn't involved in suspend, e.g., when
    we're going to the "suspend-to-idle" or "standby" system sleep states.

    Fixes: 1df5172c5c25 ("PCI: Suspend/resume quirks for Apple thunderbolt")
    Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=212767
    Link: https://lore.kernel.org/r/20210520235501.917397-1-Hi-Angel@yandex.ru
    Signed-off-by: Konstantin Kharlamov <Hi-Angel@yandex.ru>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Lukas Wunner <lukas@wunner.de>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ad365e9351ac2b450e7e79932ff6abf59342d91a
Author: Hou Tao <houtao1@huawei.com>
Date:   Thu Jun 17 15:45:47 2021 +0800

    dm btree remove: assign new_root only when removal succeeds

    commit b6e58b5466b2959f83034bead2e2e1395cca8aeb upstream.

    remove_raw() in dm_btree_remove() may fail due to IO read error
    (e.g. read the content of origin block fails during shadowing),
    and the value of shadow_spine::root is uninitialized, but
    the uninitialized value is still assign to new_root in the
    end of dm_btree_remove().

    For dm-thin, the value of pmd->details_root or pmd->root will become
    an uninitialized value, so if trying to read details_info tree again
    out-of-bound memory may occur as showed below:

      general protection fault, probably for non-canonical address 0x3fdcb14c8d7520
      CPU: 4 PID: 515 Comm: dmsetup Not tainted 5.13.0-rc6
      Hardware name: QEMU Standard PC
```

```
        RIP: 0010:metadata_ll_load_ie+0x14/0x30
        Call Trace:
         sm_metadata_count_is_more_than_one+0xb9/0xe0
         dm_tm_shadow_block+0x52/0x1c0
         shadow_step+0x59/0xf0
         remove_raw+0xb2/0x170
         dm_btree_remove+0xf4/0x1c0
         dm_pool_delete_thin_device+0xc3/0x140
         pool_message+0x218/0x2b0
         target_message+0x251/0x290
         ctl_ioctl+0x1c4/0x4d0
         dm_ctl_ioctl+0xe/0x20
         __x64_sys_ioctl+0x7b/0xb0
         do_syscall_64+0x40/0xb0
         entry_SYSCALL_64_after_hwframe+0x44/0xae

    Fixing it by only assign new_root when removal succeeds

    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3993b84d44939a0db7c549e8a25c6256420c0f4b
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Tue Jun 15 13:45:55 2021 -0400

    dm writecache: flush origin device when writing and cache is full

    commit ee55b92a7391bf871939330f662651b54be51b73 upstream.

    Commit d53f1fafec9d086f1c5166436abefdaef30e0363 ("dm writecache: do
    direct write if the cache is full") changed dm-writecache, so that it
    writes directly to the origin device if the cache is full.
    Unfortunately, it doesn't forward flush requests to the origin device,
    so that there is a bug where flushes are being ignored.

    Fix this by adding missing flush forwarding.

    For PMEM mode, we fix this bug by disabling direct writes to the origin
    device, because it performs better.

    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Fixes: d53f1fafec9d ("dm writecache: do direct write if the cache is full")
    Cc: stable@vger.kernel.org # v5.7+
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 14f19e585018870fd1fa4f237cb130d9a795d3bc
Author: Damien Le Moal <damien.lemoal@wdc.com>
Date:   Wed May 19 10:26:16 2021 +0900

    dm zoned: check zone capacity

    commit bab68499428ed934f0493ac74197ed6f36204260 upstream.

    The dm-zoned target cannot support zoned block devices with zones that
    have a capacity smaller than the zone size (e.g. NVMe zoned namespaces)
    due to the current chunk zone mapping implementation as it is assumed
    that zones and chunks have the same size with all blocks usable.
    If a zoned drive is found to have zones with a capacity different from
    the zone size, fail the target initialization.

    Signed-off-by: Damien Le Moal <damien.lemoal@wdc.com>
    Cc: stable@vger.kernel.org # v5.9+
```

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

  commit 0115687be7b13993066aef602253a53d55f5b11f
  Author: Sai Prakash Ranjan <saiprakash.ranjan@codeaurora.org>
  Date:   Mon Jun 14 11:59:00 2021 -0600

      coresight: tmc-etf: Fix global-out-of-bounds in tmc_update_etf_buffer()

      commit 5fae8a946ac2df879caf3f79a193d4766d00239b upstream.

      commit 6f755e85c332 ("coresight: Add helper for inserting synchronization
      packets") removed trailing '\0' from barrier_pkt array and updated the
      call sites like etb_update_buffer() to have proper checks for barrier_pkt
      size before read but missed updating tmc_update_etf_buffer() which still
      reads barrier_pkt past the array size resulting in KASAN out-of-bounds
      bug. Fix this by adding a check for barrier_pkt size before accessing
      like it is done in etb_update_buffer().

       BUG: KASAN: global-out-of-bounds in tmc_update_etf_buffer+0x4b8/0x698
       Read of size 4 at addr ffffffd05b7d1030 by task perf/2629

       Call trace:
        dump_backtrace+0x0/0x27c
        show_stack+0x20/0x2c
        dump_stack+0x11c/0x188
        print_address_description+0x3c/0x4a4
        __kasan_report+0x140/0x164
        kasan_report+0x10/0x18
        __asan_report_load4_noabort+0x1c/0x24
        tmc_update_etf_buffer+0x4b8/0x698
        etm_event_stop+0x248/0x2d8
        etm_event_del+0x20/0x2c
        event_sched_out+0x214/0x6f0
        group_sched_out+0xd0/0x270
        ctx_sched_out+0x2ec/0x518
        __perf_event_task_sched_out+0x4fc/0xe6c
        __schedule+0x1094/0x16a0
        preempt_schedule_irq+0x88/0x170
        arm64_preempt_schedule_irq+0xf0/0x18c
        el1_irq+0xe8/0x180
        perf_event_exec+0x4d8/0x56c
        setup_new_exec+0x204/0x400
        load_elf_binary+0x72c/0x18c0
        search_binary_handler+0x13c/0x420
        load_script+0x500/0x6c4
        search_binary_handler+0x13c/0x420
        exec_binprm+0x118/0x654
        __do_execve_file+0x77c/0xba4
        __arm64_compat_sys_execve+0x98/0xac
        el0_svc_common+0x1f8/0x5e0
        el0_svc_compat_handler+0x84/0xb0
        el0_svc_compat+0x10/0x50

       The buggy address belongs to the variable:
        barrier_pkt+0x10/0x40

       Memory state around the buggy address:
        ffffffd05b7d0f00: fa fa fa fa 04 fa fa fa fa fa fa fa 00 00 00 00
        ffffffd05b7d0f80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
       >ffffffd05b7d1000: 00 00 00 00 00 00 fa fa fa fa fa fa 00 00 00 03
                                           ^
        ffffffd05b7d1080: fa fa fa fa 00 02 fa fa fa fa fa fa 03 fa fa fa
        ffffffd05b7d1100: fa fa fa fa 00 00 00 00 05 fa fa fa fa fa fa fa
       ==================================================================
```

```
    Link: https://lore.kernel.org/r/20210505093430.18445-1-saiprakash.ranjan@codeaurora.org
    Fixes: 0c3fc4d5fa26 ("coresight: Add barrier packet for synchronisation")
    Cc: stable@vger.kernel.org
    Signed-off-by: Sai Prakash Ranjan <saiprakash.ranjan@codeaurora.org>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Link: https://lore.kernel.org/r/20210614175901.532683-6-mathieu.poirier@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2ac133f4b299bd5a76cb5be1bc49733b6d45f7a7
Author: Jeremy Linton <jeremy.linton@arm.com>
Date:   Mon Jun 14 11:59:01 2021 -0600

    coresight: Propagate symlink failure

    commit 51dd19a7e9f8fbbb7cd92b8a357091911eae7f78 upstream.

    If the symlink is unable to be created, the driver goes
    ahead and continues device creation. Instead lets propagate
    the failure, and fail the probe.

    Link: https://lore.kernel.org/r/20210526204042.2681700-1-jeremy.linton@arm.com
    Fixes: 8a7365c2d418 ("coresight: Expose device connections via sysfs")
    Cc: stable@vger.kernel.org
    Signed-off-by: Jeremy Linton <jeremy.linton@arm.com>
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Link: https://lore.kernel.org/r/20210614175901.532683-7-mathieu.poirier@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5b37fd9da03f2e60fc3f7f8e997d277f71578108
Author: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
Date:   Mon May 24 02:32:05 2021 -0700

    ipack/carriers/tpci200: Fix a double free in tpci200_pci_probe

    commit 9272e5d0028d45a3b45b58c9255e6e0df53f7ad9 upstream.

    In the out_err_bus_register error branch of tpci200_pci_probe,
    tpci200->info->cfg_regs is freed by tpci200_uninstall()->
    tpci200_unregister()->pci_iounmap(..,tpci200->info->cfg_regs)
    in the first time.

    But later, iounmap() is called to free tpci200->info->cfg_regs
    again.

    My patch sets tpci200->info->cfg_regs to NULL after tpci200_uninstall()
    to avoid the double free.

    Fixes: cea2f7cdff2af ("Staging: ipack/bridges/tpci200: Use the TPCI200 in big endian
mode")
    Cc: stable <stable@vger.kernel.org>
    Acked-by: Samuel Iglesias Gonsalvez <siglesias@igalia.com>
    Signed-off-by: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
    Link: https://lore.kernel.org/r/20210524093205.8333-1-lyl2019@mail.ustc.edu.cn
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1d2865bbd89e4397e91ef01de53da7fca89f7429
Author: Paul Burton <paulburton@google.com>
Date:   Thu Jul 1 10:24:07 2021 -0700

    tracing: Resize tgid_map to pid_max, not PID_MAX_DEFAULT

    commit 4030a6e6a6a4a42ff8c18414c9e0c93e24cc70b8 upstream.
```

```
    Currently tgid_map is sized at PID_MAX_DEFAULT entries, which means that
    on systems where pid_max is configured higher than PID_MAX_DEFAULT the
    ftrace record-tgid option doesn't work so well. Any tasks with PIDs
    higher than PID_MAX_DEFAULT are simply not recorded in tgid_map, and
    don't show up in the saved_tgids file.

    In particular since systemd v243 & above configure pid_max to its
    highest possible 1<<22 value by default on 64 bit systems this renders
    the record-tgids option of little use.

    Increase the size of tgid_map to the configured pid_max instead,
    allowing it to cover the full range of PIDs up to the maximum value of
    PID_MAX_LIMIT if the system is configured that way.

    On 64 bit systems with pid_max == PID_MAX_LIMIT this will increase the
    size of tgid_map from 256KiB to 16MiB. Whilst this 64x increase in
    memory overhead sounds significant 64 bit systems are presumably best
    placed to accommodate it, and since tgid_map is only allocated when the
    record-tgid option is actually used presumably the user would rather it
    spends sufficient memory to actually record the tgids they expect.

    The size of tgid_map could also increase for CONFIG_BASE_SMALL=y
    configurations, but these seem unlikely to be systems upon which people
    are both configuring a large pid_max and running ftrace with record-tgid
    anyway.

    Of note is that we only allocate tgid_map once, the first time that the
    record-tgid option is enabled. Therefore its size is only set once, to
    the value of pid_max at the time the record-tgid option is first
    enabled. If a user increases pid_max after that point, the saved_tgids
    file will not contain entries for any tasks with pids beyond the earlier
    value of pid_max.

    Link: https://lkml.kernel.org/r/20210701172407.889626-2-paulburton@google.com

    Fixes: d914ba37d714 ("tracing: Add support for recording tgid of tasks")
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Joel Fernandes <joelaf@google.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Paul Burton <paulburton@google.com>
    [ Fixed comment coding style ]
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 38164905b1cc30f4099d78a2394115b7046dfb0f
Author: Paul Burton <paulburton@google.com>
Date:   Tue Jun 29 17:34:05 2021 -0700

    tracing: Simplify & fix saved_tgids logic

    commit b81b3e959adb107cd5b36c7dc5ba1364bbd31eb2 upstream.

    The tgid_map array records a mapping from pid to tgid, where the index
    of an entry within the array is the pid & the value stored at that index
    is the tgid.

    The saved_tgids_next() function iterates over pointers into the tgid_map
    array & dereferences the pointers which results in the tgid, but then it
    passes that dereferenced value to trace_find_tgid() which treats it as a
    pid & does a further lookup within the tgid_map array. It seems likely
    that the intent here was to skip over entries in tgid_map for which the
    recorded tgid is zero, but instead we end up skipping over entries for
    which the thread group leader hasn't yet had its own tgid recorded in
    tgid_map.
```

```
    A minimal fix would be to remove the call to trace_find_tgid, turning:

      if (trace_find_tgid(*ptr))

    into:

      if (*ptr)

    ..but it seems like this logic can be much simpler if we simply let
    seq_read() iterate over the whole tgid_map array & filter out empty
    entries by returning SEQ_SKIP from saved_tgids_show(). Here we take that
    approach, removing the incorrect logic here entirely.

    Link: https://lkml.kernel.org/r/20210630003406.4013668-1-paulburton@google.com

    Fixes: d914ba37d714 ("tracing: Add support for recording tgid of tasks")
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Joel Fernandes <joelaf@google.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Paul Burton <paulburton@google.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 589f6fac5c926ed3e0093636db44cd5c5942d55f
Author: Jan Kara <jack@suse.cz>
Date:   Mon Jun 7 13:26:13 2021 +0200

    rq-qos: fix missed wake-ups in rq_qos_throttle try two

    commit 11c7aa0ddea8611007768d3e6b58d45dc60a19e1 upstream.

    Commit 545fbd0775ba ("rq-qos: fix missed wake-ups in rq_qos_throttle")
    tried to fix a problem that a process could be sleeping in rq_qos_wait()
    without anyone to wake it up. However the fix is not complete and the
    following can still happen:

    CPU1 (waiter1)          CPU2 (waiter2)          CPU3 (waker)
    rq_qos_wait()           rq_qos_wait()
      acquire_inflight_cb() -> fails
                               acquire_inflight_cb() -> fails

                                                     completes IOs, inflight
                                                       decreased
      prepare_to_wait_exclusive()
                               prepare_to_wait_exclusive()
      has_sleeper = !wq_has_single_sleeper() -> true as there are two sleepers
                               has_sleeper = !wq_has_single_sleeper() -> true
      io_schedule()            io_schedule()

    Deadlock as now there's nobody to wakeup the two waiters. The logic
    automatically blocking when there are already sleepers is really subtle
    and the only way to make it work reliably is that we check whether there
    are some waiters in the queue when adding ourselves there. That way, we
    are guaranteed that at least the first process to enter the wait queue
    will recheck the waiting condition before going to sleep and thus
    guarantee forward progress.

    Fixes: 545fbd0775ba ("rq-qos: fix missed wake-ups in rq_qos_throttle")
    CC: stable@vger.kernel.org
    Signed-off-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20210607112613.25344-1-jack@suse.cz
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d57fcab190b60f43046d5836c3c56114b4f50080
```

```
Author: Yun Zhou <yun.zhou@windriver.com>
Date:   Sat Jun 26 11:21:55 2021 +0800

    seq_buf: Fix overflow in seq_buf_putmem_hex()

    commit d3b16034a24a112bb83aeb669ac5b9b01f744bb7 upstream.

    There's two variables being increased in that loop (i and j), and i
    follows the raw data, and j follows what is being written into the buffer.
    We should compare 'i' to MAX_MEMHEX_BYTES or compare 'j' to HEX_CHARS.
    Otherwise, if 'j' goes bigger than HEX_CHARS, it will overflow the
    destination buffer.

    Link: https://lore.kernel.org/lkml/20210625122453.5e2fe304@oasis.local.home/
    Link: https://lkml.kernel.org/r/20210626032156.47889-1-yun.zhou@windriver.com

    Cc: stable@vger.kernel.org
    Fixes: 5e3ca0ec76fce ("ftrace: introduce the "hex" output method")
    Signed-off-by: Yun Zhou <yun.zhou@windriver.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 775ab9f4a4d9027ea55e9cd683e9ff48d09d749a
Author: Kees Cook <keescook@chromium.org>
Date:   Wed Jun 23 13:39:33 2021 -0700

    lkdtm: Enable DOUBLE_FAULT on all architectures

    commit f123c42bbeff26bfe8bdb08a01307e92d51eec39 upstream.

    Where feasible, I prefer to have all tests visible on all architectures,
    but to have them wired to XFAIL. DOUBLE_FAIL was set up to XFAIL, but
    wasn't actually being added to the test list.

    Fixes: cea23efb4de2 ("lkdtm/bugs: Make double-fault test always available")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210623203936.3151093-7-keescook@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 42e737661a0dfd872c1392a33b20273127a54d27
Author: Ferry Toth <ftoth@exalondelft.nl>
Date:   Tue May 18 23:27:09 2021 +0200

    extcon: intel-mrfld: Sync hardware and software state on init

    commit ecb5bdff901139850fb3ca3ae2d0cccac045bc52 upstream.

    extcon driver for Basin Cove PMIC shadows the switch status used for dwc3
    DRD to detect a change in the switch position. This change initializes the
    status at probe time.

    Cc: stable@vger.kernel.org
    Fixes: 492929c54791 ("extcon: mrfld: Introduce extcon driver for Basin Cove PMIC")
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Ferry Toth <ftoth@exalondelft.nl>
    Signed-off-by: Chanwoo Choi <cw00.choi@samsung.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6fd650c47411ee353b1efba1500c67b170d7d0c9
Author: Kees Cook <keescook@chromium.org>
Date:   Wed Jun 23 13:39:29 2021 -0700

    selftests/lkdtm: Fix expected text for CR4 pinning
```

```
    commit c2eb472bbe25b3f360990f23b293b3fbadfa4bc0 upstream.

    The error text for CR4 pinning changed. Update the test to match.

    Fixes: a13b9d0b9721 ("x86/cpu: Use pinning mask for CR4 bits needing to be 0")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210623203936.3151093-3-keescook@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 07b761cf78b147b035eb8574cc22bd016e51b307
Author: Kees Cook <keescook@chromium.org>
Date:   Wed Jun 23 13:39:31 2021 -0700

    lkdtm/bugs: XFAIL UNALIGNED_LOAD_STORE_WRITE

    commit a15676ac8f24a9ac5fd881cf17be4be13fa0910a upstream.

    When built under CONFIG_HAVE_EFFICIENT_UNALIGNED_ACCESS, this test is
    expected to fail (i.e. not trip an exception).

    Fixes: 46d1a0f03d66 ("selftests/lkdtm: Add tests for LKDTM targets")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210623203936.3151093-5-keescook@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 866573b6c9bfc1feb9d8827a64c17ad21845a6b0
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Fri Jun 11 11:23:21 2021 +0100

    nvmem: core: add a missing of_node_put

    commit 63879e2964bceee2aa5bbe8b99ea58bba28bb64f upstream.

    'for_each_child_of_node' performs an of_node_get on each iteration, so a
    return from the middle of the loop requires an of_node_put.

    Fixes: e888d445ac33 ("nvmem: resolve cells from DT at registration time")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20210611102321.11509-1-srinivas.kandagatla@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9c6c73fc99b9943b9420534dc47130c312fc1eaa
Author: Limeng <Meng.Li@windriver.com>
Date:   Wed Apr 7 13:25:25 2021 +0800

    mfd: syscon: Free the allocated name field of struct regmap_config

    commit 56a1188159cb2b87fbcb5a7a7afb38a4dd9db0c1 upstream.

    The commit 529a1101212a("mfd: syscon: Don't free allocated name
    for regmap_config") doesn't free the allocated name field of struct
    regmap_config, but introduce a memory leak. There is another
    commit 94cc89eb8fa5("regmap: debugfs: Fix handling of name string
    for debugfs init delays") fixing this debugfs init issue from root
    cause. With this fixing, the name field in struct regmap_debugfs_node
    is removed. When initialize debugfs for syscon driver, the name
    field of struct regmap_config is not used anymore. So, the allocated
    name field of struct regmap_config is need to be freed directly after
    regmap initialization to avoid memory leak.

    Cc: stable@vger.kernel.org
```

```
    Fixes: 529a1101212a("mfd: syscon: Don't free allocated name for regmap_config")
    Signed-off-by: Meng Li <Meng.Li@windriver.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8451754bca7ba463838627969653e1ea94c26f72
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun Jun 27 01:47:49 2021 +0200

    power: supply: ab8500: Fix an old bug

    commit f1c74a6c07e76fcb31a4bcc1f437c4361a2674ce upstream.

    Trying to get the AB8500 charging driver working I ran into a bit
    of bitrot: we haven't used the driver for a while so errors in
    refactorings won't be noticed.

    This one is pretty self evident: use argument to the macro or we
    end up with a random pointer to something else.

    Cc: stable@vger.kernel.org
    Cc: Krzysztof Kozlowski <krzk@kernel.org>
    Cc: Marcus Cooper <codekipper@gmail.com>
    Fixes: 297d716f6260 ("power_supply: Change ownership from driver to core")
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c0756f75c22149d20fcb7d8409827cee905eb386
Author: Zhihao Cheng <chengzhihao1@huawei.com>
Date:   Mon May 31 20:52:09 2021 +0800

    ubifs: Fix races between xattr_{set|get} and listxattr operations

    commit f4e3634a3b642225a530c292fdb1e8a4007507f5 upstream.

    UBIFS may occur some problems with concurrent xattr_{set|get} and
    listxattr operations, such as assertion failure, memory corruption,
    stale xattr value[1].

    Fix it by importing a new rw-lock in @ubifs_inode to serilize write
    operations on xattr, concurrent read operations are still effective,
    just like ext4.

    [1] https://lore.kernel.org/linux-mtd/20200630130438.141649-1-houtao1@huawei.com

    Fixes: 1e51764a3c2ac05a23 ("UBIFS: add new flash file system")
    Cc: stable@vger.kernel.org  # v2.6+
    Signed-off-by: Zhihao Cheng <chengzhihao1@huawei.com>
    Reviewed-by: Sascha Hauer <s.hauer@pengutronix.de>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec00b71f402aa4305e8df943869bf2c75adb995a
Author: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
Date:   Mon Jun 28 14:58:03 2021 -0700

    thermal/drivers/int340x/processor_thermal: Fix tcc setting

    commit fe6a6de6692e7f7159c1ff42b07ecd737df712b4 upstream.

    The following fixes are done for tcc sysfs interface:
    - TCC is 6 bits only from bit 29-24
    - TCC of 0 is valid
    - When BIT(31) is set, this register is read only
```

```
      - Check for invalid tcc value
      - Error for negative values

      Fixes: fdf4f2fb8e899 ("drivers: thermal: processor_thermal_device: Export sysfs
  interface for TCC offset")
      Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
      Cc: stable@vger.kernel.org
      Acked-by: Zhang Rui <rui.zhang@intel.com>
      Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
      Link: https://lore.kernel.org/r/20210628215803.75038-1-
  srinivas.pandruvada@linux.intel.com
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 238a7c0a5c3090b46e68de10003df8aebe5d8935
  Author: Petr Pavlu <petr.pavlu@suse.com>
  Date:   Thu May 13 14:26:36 2021 +0200

      ipmi/watchdog: Stop watchdog timer when the current action is 'none'

      commit 2253042d86f57d90a621ac2513a7a7a13afcf809 upstream.

      When an IPMI watchdog timer is being stopped in ipmi_close() or
      ipmi_ioctl(WDIOS_DISABLECARD), the current watchdog action is updated to
      WDOG_TIMEOUT_NONE and _ipmi_set_timeout(IPMI_SET_TIMEOUT_NO_HB) is called
      to install this action. The latter function ends up invoking
      __ipmi_set_timeout() which makes the actual 'Set Watchdog Timer' IPMI
      request.

      For IPMI 1.0, this operation results in fully stopping the watchdog timer.
      For IPMI >= 1.5, function __ipmi_set_timeout() always specifies the "don't
      stop" flag in the prepared 'Set Watchdog Timer' IPMI request. This causes
      that the watchdog timer has its action correctly updated to 'none' but the
      timer continues to run. A problem is that IPMI firmware can then still log
      an expiration event when the configured timeout is reached, which is
      unexpected because the watchdog timer was requested to be stopped.

      The patch fixes this problem by not setting the "don't stop" flag in
      __ipmi_set_timeout() when the current action is WDOG_TIMEOUT_NONE which
      results in stopping the watchdog timer. This makes the behaviour for
      IPMI >= 1.5 consistent with IPMI 1.0. It also matches the logic in
      __ipmi_heartbeat() which does not allow to reset the watchdog if the
      current action is WDOG_TIMEOUT_NONE as that would start the timer.

      Signed-off-by: Petr Pavlu <petr.pavlu@suse.com>
      Message-Id: <10a41bdc-9c99-089c-8d89-fa98ce5ea080@suse.com>
      Cc: stable@vger.kernel.org
      Signed-off-by: Corey Minyard <cminyard@mvista.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 13b307880b66ef6e7f4eeee2a26cbbe0f8206863
  Author: Nathan Chancellor <nathan@kernel.org>
  Date:   Thu Feb 11 12:42:58 2021 -0700

      qemu_fw_cfg: Make fw_cfg_rev_attr a proper kobj_attribute

      commit fca41af18e10318e4de090db47d9fa7169e1bf2f upstream.

      fw_cfg_showrev() is called by an indirect call in kobj_attr_show(),
      which violates clang's CFI checking because fw_cfg_showrev()'s second
      parameter is 'struct attribute', whereas the ->show() member of 'struct
      kobj_structure' expects the second parameter to be of type 'struct
      kobj_attribute'.

      $ cat /sys/firmware/qemu_fw_cfg/rev
      3
```

```
    $ dmesg | grep "CFI failure"
    [   26.016832] CFI failure (target: fw_cfg_showrev+0x0/0x8):

    Fix this by converting fw_cfg_rev_attr to 'struct kobj_attribute' where
    this would have been caught automatically by the incompatible pointer
    types compiler warning. Update fw_cfg_showrev() accordingly.

    Fixes: 75f3e8e47f38 ("firmware: introduce sysfs driver for QEMU's fw_cfg device")
    Link: https://github.com/ClangBuiltLinux/linux/issues/1299
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Sami Tolvanen <samitolvanen@google.com>
    Tested-by: Sedat Dilek <sedat.dilek@gmail.com>
    Reviewed-by: Sami Tolvanen <samitolvanen@google.com>
    Reviewed-by: Philippe Mathieu-Daudé <philmd@redhat.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20210211194258.4137998-1-nathan@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e71e1344d61509a5b444b1a4cf0c0ea19fd9dd8
Author: Jesse Brandeburg <jesse.brandeburg@intel.com>
Date:   Fri May 7 11:56:25 2021 -0700

    i40e: fix PTP on 5Gb links

    commit 26b0ce8dd3dd704393dbace4dc416adfeffe531f upstream.

    As reported by Alex Sergeev, the i40e driver is incrementing the PTP
    clock at 40Gb speeds when linked at 5Gb. Fix this bug by making
    sure that the right multiplier is selected when linked at 5Gb.

    Fixes: 3dbdd6c2f70a ("i40e: Add support for 5Gbps cards")
    Cc: stable@vger.kernel.org
    Reported-by: Alex Sergeev <asergeev@carbonrobotics.com>
    Suggested-by: Alex Sergeev <asergeev@carbonrobotics.com>
    Signed-off-by: Jesse Brandeburg <jesse.brandeburg@intel.com>
    Tested-by: Tony Brelinski <tonyx.brelinski@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 35af69c7c0490fdccfc159c6a87e4d1dc070838a
Author: Brian Norris <briannorris@chromium.org>
Date:   Fri May 14 19:42:27 2021 -0700

    mwifiex: bring down link before deleting interface

    commit 1f9482aa8d412b4ba06ce6ab8e333fb8ca29a06e upstream.

    We can deadlock when rmmod'ing the driver or going through firmware
    reset, because the cfg80211_unregister_wdev() has to bring down the link
    for us, ... which then grab the same wiphy lock.

    nl80211_del_interface() already handles a very similar case, with a nice
    description:

            /*
             * We hold RTNL, so this is safe, without RTNL opencount cannot
             * reach 0, and thus the rdev cannot be deleted.
             *
             * We need to do it for the dev_close(), since that will call
             * the netdev notifiers, and we need to acquire the mutex there
             * but don't know if we get there from here or from some other
             * place (e.g. "ip link set ... down").
             */
```

```
            mutex_unlock(&rdev->wiphy.mtx);
    ...

    Do similarly for mwifiex teardown, by ensuring we bring the link down
    first.

    Sample deadlock trace:

    [  247.103516] INFO: task rmmod:2119 blocked for more than 123 seconds.
    [  247.110630]       Not tainted 5.12.4 #5
    [  247.115796] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
    [  247.124557] task:rmmod           state:D stack:    0 pid: 2119 ppid:  2114
 flags:0x00400208
    [  247.133905] Call trace:
    [  247.136644]  __switch_to+0x130/0x170
    [  247.140643]  __schedule+0x714/0xa0c
    [  247.144548]  schedule_preempt_disabled+0x88/0xf4
    [  247.149714]  __mutex_lock_common+0x43c/0x750
    [  247.154496]  mutex_lock_nested+0x5c/0x68
    [  247.158884]  cfg80211_netdev_notifier_call+0x280/0x4e0 [cfg80211]
    [  247.165769]  raw_notifier_call_chain+0x4c/0x78
    [  247.170742]  call_netdevice_notifiers_info+0x68/0xa4
    [  247.176305]  __dev_close_many+0x7c/0x138
    [  247.180693]  dev_close_many+0x7c/0x10c
    [  247.184893]  unregister_netdevice_many+0xfc/0x654
    [  247.190158]  unregister_netdevice_queue+0xb4/0xe0
    [  247.195424]  _cfg80211_unregister_wdev+0xa4/0x204 [cfg80211]
    [  247.201816]  cfg80211_unregister_wdev+0x20/0x2c [cfg80211]
    [  247.208016]  mwifiex_del_virtual_intf+0xc8/0x188 [mwifiex]
    [  247.214174]  mwifiex_uninit_sw+0x158/0x1b0 [mwifiex]
    [  247.219747]  mwifiex_remove_card+0x38/0xa0 [mwifiex]
    [  247.225316]  mwifiex_pcie_remove+0xd0/0xe0 [mwifiex_pcie]
    [  247.231451]  pci_device_remove+0x50/0xe0
    [  247.235849]  device_release_driver_internal+0x110/0x1b0
    [  247.241701]  driver_detach+0x5c/0x9c
    [  247.245704]  bus_remove_driver+0x84/0xb8
    [  247.250095]  driver_unregister+0x3c/0x60
    [  247.254486]  pci_unregister_driver+0x2c/0x90
    [  247.259267]  cleanup_module+0x18/0xcdc [mwifiex_pcie]

    Fixes: a05829a7222e ("cfg80211: avoid holding the RTNL when calling the driver")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/linux-wireless/98392296-40ee-6300-369c-
 32e16cff3725@gmail.com/
    Link: https://lore.kernel.org/linux-wireless/ab4d00ce52f32bd8e45ad0448a44737e@bewaar.me/
    Reported-by: Maximilian Luz <luzmaximilian@gmail.com>
    Reported-by: dave@bewaar.me
    Cc: Johannes Berg <johannes@sipsolutions.net>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Tested-by: Maximilian Luz <luzmaximilian@gmail.com>
    Tested-by: Dave Olsthoorn <dave@bewaar.me>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210515024227.2159311-1-briannorris@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 5665e64c99e35c203525f2645df08bc5766f7a20
 Author: Dmitry Osipenko <digetx@gmail.com>
 Date:   Sat May 29 18:46:46 2021 +0300

    ASoC: tegra: Set driver_name=tegra for all machine drivers

    commit f6eb84fa596abf28959fc7e0b626f925eb1196c7 upstream.

    The driver_name="tegra" is now required by the newer ALSA UCMs, otherwise
    Tegra UCMs don't match by the path/name.
```

```
    All Tegra machine drivers are specifying the card's name, but it has no
    effect if model name is specified in the device-tree since it overrides
    the card's name. We need to set the driver_name to "tegra" in order to
    get a usable lookup path for the updated ALSA UCMs. The new UCM lookup
    path has a form of driver_name/card_name.

    The old lookup paths that are based on driver module name continue to
    work as before. Note that UCM matching never worked for Tegra ASoC drivers
    if they were compiled as built-in, this is fixed by supporting the new
    naming scheme.

    Cc: stable@vger.kernel.org
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Link: https://lore.kernel.org/r/20210529154649.25936-2-digetx@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 72104552eae2cf6979129e90b2fbb39951e24496
Author: Russ Weight <russell.h.weight@intel.com>
Date:   Mon Jun 14 10:09:03 2021 -0700

    fpga: stratix10-soc: Add missing fpga_mgr_free() call

    commit d9ec9daa20eb8de1efe6abae78c9835ec8ed86f9 upstream.

    The stratix10-soc driver uses fpga_mgr_create() function and is therefore
    responsible to call fpga_mgr_free() to release the class driver resources.
    Add a missing call to fpga_mgr_free in the s10_remove() function.

    Signed-off-by: Russ Weight <russell.h.weight@intel.com>
    Reviewed-by: Xu Yilun <yilun.xu@intel.com>
    Signed-off-by: Moritz Fischer <mdf@kernel.org>
    Fixes: e7eef1d7633a ("fpga: add intel stratix10 soc fpga manager driver")
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210614170909.232415-3-mdf@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c3f4ad1ba2eb08ef02df8adbe3827d3715003436
Author: Samuel Holland <samuel@sholland.org>
Date:   Fri May 14 21:14:39 2021 -0500

    clocksource/arm_arch_timer: Improve Allwinner A64 timer workaround

    commit 8b33dfe0ba1c84c1aab2456590b38195837f1e6e upstream.

    Bad counter reads are experienced sometimes when bit 10 or greater rolls
    over. Originally, testing showed that at least 10 lower bits would be
    set to the same value during these bad reads. However, some users still
    reported time skips.

    Wider testing revealed that on some chips, occasionally only the lowest
    9 bits would read as the anomalous value. During these reads (which
    still happen only when bit 10), bit 9 would read as the correct value.

    Reduce the mask by one bit to cover these cases as well.

    Cc: stable@vger.kernel.org
    Fixes: c950ca8c35ee ("clocksource/drivers/arch_timer: Workaround for Allwinner A64 timer
instability")
    Reported-by: Roman Stratiienko <r.stratiienko@gmail.com>
    Signed-off-by: Samuel Holland <samuel@sholland.org>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20210515021439.55316-1-samuel@sholland.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit a34295528a14a733f38b59f6f324d6e34d61d21f
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Sat Mar 27 22:01:36 2021 +0100

    cpu/hotplug: Cure the cpusets trainwreck

    commit b22afcdf04c96ca58327784e280e10288cfd3303 upstream.

    Alexey and Joshua tried to solve a cpusets related hotplug problem which is
    user space visible and results in unexpected behaviour for some time after
    a CPU has been plugged in and the corresponding uevent was delivered.

    cpusets delegate the hotplug work (rebuilding cpumasks etc.) to a
    workqueue. This is done because the cpusets code has already a lock
    nesting of cgroups_mutex -> cpu_hotplug_lock. A synchronous callback or
    waiting for the work to finish with cpu_hotplug_lock held can and will
    deadlock because that results in the reverse lock order.

    As a consequence the uevent can be delivered before cpusets have consistent
    state which means that a user space invocation of sched_setaffinity() to
    move a task to the plugged CPU fails up to the point where the scheduled
    work has been processed.

    The same is true for CPU unplug, but that does not create user observable
    failure (yet).

    It's still inconsistent to claim that an operation is finished before it
    actually is and that's the real issue at hand. uevents just make it
    reliably observable.

    Obviously the problem should be fixed in cpusets/cgroups, but untangling
    that is pretty much impossible because according to the changelog of the
    commit which introduced this 8 years ago:

     3a5a6d0c2b03("cpuset: don't nest cgroup_mutex inside get_online_cpus()")

    the lock order cgroups_mutex -> cpu_hotplug_lock is a design decision and
    the whole code is built around that.

    So bite the bullet and invoke the relevant cpuset function, which waits for
    the work to finish, in _cpu_up/down() after dropping cpu_hotplug_lock and
    only when tasks are not frozen by suspend/hibernate because that would
    obviously wait forever.

    Waiting there with cpu_add_remove_lock, which is protecting the present
    and possible CPU maps, held is not a problem at all because neither work
    queues nor cpusets/cgroups have any lockchains related to that lock.

    Waiting in the hotplug machinery is not problematic either because there
    are already state callbacks which wait for hardware queues to drain. It
    makes the operations slightly slower, but hotplug is slow anyway.

    This ensures that state is consistent before returning from a hotplug
    up/down operation. It's still inconsistent during the operation, but that's
    a different story.

    Add a large comment which explains why this is done and why this is not a
    dump ground for the hack of the day to work around half thought out locking
    schemes. Document also the implications vs. hotplug operations and
    serialization or the lack of it.

    Thanks to Alexy and Joshua for analyzing why this temporary
    sched_setaffinity() failure happened.
```

```
    Fixes: 3a5a6d0c2b03("cpuset: don't nest cgroup_mutex inside get_online_cpus()")
    Reported-by: Alexey Klimov <aklimov@redhat.com>
    Reported-by: Joshua Baker <jobaker@redhat.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Alexey Klimov <aklimov@redhat.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/87tuowcnv3.ffs@nanos.tec.linutronix.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dcb4b49e3ced01bb3e3ac69cb394accf5615cda1
Author: Zhenyu Ye <yezhenyu2@huawei.com>
Date:   Wed Jun 23 15:05:22 2021 +0800

    arm64: tlb: fix the TTL value of tlb_get_level

    commit 52218fcd61cb42bde0d301db4acb3ffdf3463cc7 upstream.

    The TTL field indicates the level of page table walk holding the *leaf*
    entry for the address being invalidated. But currently, the TTL field
    may be set to an incorrent value in the following stack:

    pte_free_tlb
        __pte_free_tlb
            tlb_remove_table
                tlb_table_invalidate
                    tlb_flush_mmu_tlbonly
                        tlb_flush

    In this case, we just want to flush a PTE page, but the tlb->cleared_pmds
    is set and we get tlb_level = 2 in the tlb_get_level() function. This may
    cause some unexpected problems.

    This patch set the TTL field to 0 if tlb->freed_tables is set. The
    tlb->freed_tables indicates page table pages are freed, not the leaf
    entry.

    Cc: <stable@vger.kernel.org> # 5.9.x
    Fixes: c4ab2cbc1d87 ("arm64: tlb: Set the TTL field in flush_tlb_range")
    Acked-by: Catalin Marinas <catalin.marinas@arm.com>
    Reported-by: ZhuRui <zhurui3@huawei.com>
    Signed-off-by: Zhenyu Ye <yezhenyu2@huawei.com>
    Link: https://lore.kernel.org/r/b80ead47-1f88-3a00-18e1-cacc22f54cc4@huawei.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 31d1d7db49785d74b388c55a5e33bad02061ef35
Author: Timo Sigurdsson <public_timo.s@silentcreek.de>
Date:   Mon Jun 14 09:25:39 2021 +0200

    ata: ahci_sunxi: Disable DIPM

    commit f6bca4d91b2ea052e917cca3f9d866b5cc1d500a upstream.

    DIPM is unsupported or broken on sunxi. Trying to enable the power
    management policy med_power_with_dipm on an Allwinner A20 SoC based board
    leads to immediate I/O errors and the attached SATA disk disappears from
    the /dev filesystem. A reset (power cycle) is required to make the SATA
    controller or disk work again. The A10 and A20 SoC data sheets and manuals
    don't mention DIPM at all [1], so it's fair to assume that it's simply not
    supported. But even if it was, it should be considered broken and best be
    disabled in the ahci_sunxi driver.

    [1] https://github.com/allwinner-zh/documents/tree/master/

    Fixes: c5754b5220f0 ("ARM: sunxi: Add support for Allwinner SUNXi SoCs sata to
```

```
ahci_platform")
    Cc: stable@vger.kernel.org
    Signed-off-by: Timo Sigurdsson <public_timo.s@silentcreek.de>
    Tested-by: Timo Sigurdsson <public_timo.s@silentcreek.de>
    Link: https://lore.kernel.org/r/20210614072539.3307-1-public_timo.s@silentcreek.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b7c6122f39028cffbd48b67e76ebdc0a12deed4
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Jun 17 15:58:08 2021 -0700

    docs: Makefile: Use CONFIG_SHELL not SHELL

    commit 222a28edce38b62074a950fb243df621c602b4d3 upstream.

    Fix think-o about which variable to find the Kbuild-configured shell.
    This has accidentally worked due to most shells setting $SHELL by
    default.

    Fixes: 51e46c7a4007 ("docs, parallelism: Rearrange how jobserver reservations are made")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210617225808.3907377-1-keescook@chromium.org
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5c8e4b20e70f948afb618742de76cd93e959458a
Author: Christian Löhle <CLoehle@hyperstone.com>
Date:   Wed May 12 16:03:24 2021 +0000

    mmc: core: Allow UHS-I voltage switch for SDSC cards if supported

    commit 09247e110b2efce3a104e57e887c373e0a57a412 upstream.

    While initializing an UHS-I SD card, the mmc core first tries to switch to
    1.8V I/O voltage, before it continues to change the settings for the bus
    speed mode.

    However, the current behaviour in the mmc core is inconsistent and doesn't
    conform to the SD spec. More precisely, an SD card that supports UHS-I must
    set both the SD_OCR_CCS bit and the SD_OCR_S18R bit in the OCR register
    response. When switching to 1.8V I/O the mmc core correctly checks both of
    the bits, but only the SD_OCR_S18R bit when changing the settings for bus
    speed mode.

    Rather than actually fixing the code to confirm to the SD spec, let's
    deliberately deviate from it by requiring only the SD_OCR_S18R bit for both
    parts. This enables us to support UHS-I for SDSC cards (outside spec),
    which is actually being supported by some existing SDSC cards. Moreover,
    this fixes the inconsistent behaviour.

    Signed-off-by: Christian Loehle <cloehle@hyperstone.com>
    Link:
https://lore.kernel.org/r/CWXP265MB26803AE79E0AD5ED083BF2A6C4529@CWXP265MB2680.GBRP265.PROD.
OUTLOOK.COM
    Cc: stable@vger.kernel.org
    [Ulf: Rewrote commit message and comments to clarify the changes]
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3d2e0431e68780cb90996dcedf1c1acc27a87005
Author: Wolfram Sang <wsa+renesas@sang-engineering.com>
Date:   Thu Jun 24 17:16:14 2021 +0200
```

```
    mmc: core: clear flags before allowing to retune

    commit 77347eda64ed5c9383961d1de9165f9d0b7d8df6 upstream.

    It might be that something goes wrong during tuning so the MMC core will
    immediately trigger a retune. In our case it was:

     - we sent a tuning block
     - there was an error so we need to send an abort cmd to the eMMC
     - the abort cmd had a CRC error
     - retune was set by the MMC core

    This lead to a vicious circle causing a performance regression of 75%.
    So, clear retuning flags before we enable retuning to start with a known
    cleared state.

    Reported-by Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Suggested-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Reviewed-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Tested-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Fixes: bd11e8bd03ca ("mmc: core: Flag re-tuning is needed on CRC errors")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20210624151616.38770-2-wsa+renesas@sang-engineering.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c1b907e43b320c363458ef12efbed1675c2753ea
Author: Al Cooper <alcooperx@gmail.com>
Date:   Thu Jun 24 12:30:45 2021 -0400

    mmc: sdhci: Fix warning message when accessing RPMB in HS400 mode

    commit d0244847f9fc5e20df8b7483c8a4717fe0432d38 upstream.

    When an eMMC device is being run in HS400 mode, any access to the
    RPMB device will cause the error message "mmc1: Invalid UHS-I mode
    selected". This happens as a result of tuning being disabled before
    RPMB access and then re-enabled after the RPMB access is complete.
    When tuning is re-enabled, the system has to switch from HS400
    to HS200 to do the tuning and then back to HS400. As part of
    sequence to switch from HS400 to HS200 the system is temporarily
    put into HS mode. When switching to HS mode, sdhci_get_preset_value()
    is called and does not have support for HS mode and prints the warning
    message and returns the preset for SDR12. The fix is to add support
    for MMC and SD HS modes to sdhci_get_preset_value().

    This can be reproduced on any system running eMMC in HS400 mode
    (not HS400ES) by using the "mmc" utility to run the following
    command: "mmc rpmb read-counter /dev/mmcblk0rpmb".

    Signed-off-by: Al Cooper <alcooperx@gmail.com>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Fixes: 52983382c74f ("mmc: sdhci: enhance preset value function")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20210624163045.33651-1-alcooperx@gmail.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7210a1b23c50517031c88bb8dd41bcc3c14b7a47
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon May 3 11:21:57 2021 +0200

    mmc: sdhci-acpi: Disable write protect detection on Toshiba Encore 2 WT8-B
```

```
    commit 94ee6782e045645abd9180ab9369b01293d862bd upstream.

    On the Toshiba Encore 2 WT8-B the  microSD slot always reports the card
    being write-protected even though microSD cards do not have a write-protect
    switch at all.

    Add a new DMI_QUIRK_SD_NO_WRITE_PROTECT quirk entry to sdhci-acpi.c's
    DMI quirk table for this.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Link: https://lore.kernel.org/r/20210503092157.5689-1-hdegoede@redhat.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 970aa4aefaab2d7fcceeccc3ff6b3f40b66f36e7
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Jun 17 14:33:01 2021 -0700

    drm/i915/display: Do not zero past infoframes.vsc

    commit 07b72960d2b4a087ff2445e286159e69742069cc upstream.

    intel_dp_vsc_sdp_unpack() was using a memset() size (36, struct dp_sdp)
    larger than the destination (24, struct drm_dp_vsc_sdp), clobbering
    fields in struct intel_crtc_state after infoframes.vsc. Use the actual
    target size for the memset().

    Fixes: 1b404b7dbb10 ("drm/i915/dp: Read out DP SDPs")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: José Roberto de Souza <jose.souza@intel.com>
    Signed-off-by: José Roberto de Souza <jose.souza@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210617213301.1824728-1-
keescook@chromium.org
    (cherry picked from commit c88e2647c5bb45d04dc4302018ebe6ebbf331823)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b4772d4e8ec373be3dcc931fa29516cd4c0f451d
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Mon Mar 29 18:50:45 2021 +0100

    drm/ingenic: Switch IPU plane to type OVERLAY

    commit 68b433fe6937cfa3f8975d18643d5956254edd6a upstream.

    It should have been an OVERLAY from the beginning. The documentation
    stipulates that there should be an unique PRIMARY plane per CRTC.

    Fixes: fc1acf317b01 ("drm/ingenic: Add support for the IPU")
    Cc: <stable@vger.kernel.org> # 5.8+
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Acked-by: Simon Ser <contact@emersion.fr>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210329175046.214629-2-
paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0256fdf3b07c96a77c701b11663661459f8f1dbb
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Tue Apr 27 11:20:16 2021 +0200

    drm/nouveau: Don't set allow_fb_modifiers explicitly
```

```
    commit cee93c028288b9af02919f3bd8593ba61d1e610d upstream.

    Since

    commit 890880ddfdbe256083170866e49c87618b706ac7
    Author: Paul Kocialkowski <paul.kocialkowski@bootlin.com>
    Date:   Fri Jan 4 09:56:10 2019 +0100

        drm: Auto-set allow_fb_modifiers when given modifiers at plane init

    this is done automatically as part of plane init, if drivers set the
    modifier list correctly. Which is the case here.

    Note that this fixes an inconsistency: We've set the cap everywhere,
    but only nv50+ supports modifiers. Hence cc stable, but not further
    back then the patch from Paul.

    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Cc: stable@vger.kernel.org # v5.1 +
    Cc: Pekka Paalanen <pekka.paalanen@collabora.com>
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Cc: Ben Skeggs <bskeggs@redhat.com>
    Cc: nouveau@lists.freedesktop.org
    Link: https://patchwork.freedesktop.org/patch/msgid/20210427092018.832258-6-
daniel.vetter@ffwll.ch
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c5960e823bf88beedd24ce4d9781fc1152605d3
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Tue Apr 27 11:20:12 2021 +0200

    drm/arm/malidp: Always list modifiers

    commit 26c3e7fd5a3499e408915dadae5d5360790aae9a upstream.

    Even when all we support is linear, make that explicit. Otherwise the
    uapi is rather confusing.

    Acked-by: Liviu Dudau <liviu.dudau@arm.com>
    Acked-by: Pekka Paalanen <pekka.paalanen@collabora.com>
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Cc: stable@vger.kernel.org
    Cc: Pekka Paalanen <pekka.paalanen@collabora.com>
    Cc: Liviu Dudau <liviu.dudau@arm.com>
    Cc: Brian Starkey <brian.starkey@arm.com>
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210427092018.832258-2-
daniel.vetter@ffwll.ch
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2c67909bb71b41462cb3aa251dc0c3678f27bd6
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Tue Apr 27 11:20:15 2021 +0200

    drm/msm/mdp4: Fix modifier support enabling

    commit 35cbb8c91e9cf310277d3dfb4d046df8edf2df33 upstream.

    Setting the cap without the modifier list is very confusing to
    userspace. Fix that by listing the ones we support explicitly.

    Stable backport so that userspace can rely on this working in a
    reasonable way, i.e. that the cap set implies IN_FORMATS is available.
```

```
    Acked-by: Pekka Paalanen <pekka.paalanen@collabora.com>
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Cc: stable@vger.kernel.org
    Cc: Pekka Paalanen <pekka.paalanen@collabora.com>
    Cc: Rob Clark <robdclark@chromium.org>
    Cc: Jordan Crouse <jordan@cosmicpenguin.net>
    Cc: Emil Velikov <emil.velikov@collabora.com>
    Cc: Sam Ravnborg <sam@ravnborg.org>
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210427092018.832258-5-
daniel.vetter@ffwll.ch
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 337e397a1e9e7da725505d7b4021b90d10792aaa
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Tue Apr 13 11:49:01 2021 +0200

    drm/tegra: Don't set allow_fb_modifiers explicitly

    commit be4306ad928fcf736cbe2616b6dd19d91f1bc083 upstream.

    Since

    commit 890880ddfdbe256083170866e49c87618b706ac7
    Author: Paul Kocialkowski <paul.kocialkowski@bootlin.com>
    Date:   Fri Jan 4 09:56:10 2019 +0100

        drm: Auto-set allow_fb_modifiers when given modifiers at plane init

    this is done automatically as part of plane init, if drivers set the
    modifier list correctly. Which is the case here.

    It was slightly inconsistently though, since planes with only linear
    modifier support haven't listed that explicitly. Fix that, and cc:
    stable to allow userspace to rely on this. Again don't backport
    further than where Paul's patch got added.

    Cc: stable@vger.kernel.org # v5.1 +
    Cc: Pekka Paalanen <pekka.paalanen@collabora.com>
    Acked-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Cc: Thierry Reding <thierry.reding@gmail.com>
    Cc: Jonathan Hunter <jonathanh@nvidia.com>
    Cc: linux-tegra@vger.kernel.org
    Link: https://patchwork.freedesktop.org/patch/msgid/20210413094904.3736372-10-
daniel.vetter@ffwll.ch
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5c2196d8286e1ba3a5c0a0c1466dc672530e7a3f
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Tue Mar 23 14:40:08 2021 +0000

    drm/ingenic: Fix pixclock rate for 24-bit serial panels

    commit 60a6b73dd821e98fe958b2a83393ccd724b306b1 upstream.

    When using a 24-bit panel on a 8-bit serial bus, the pixel clock
    requested by the panel has to be multiplied by 3, since the subpixels
    are shifted sequentially.

    The code (in ingenic_drm_encoder_atomic_check) already computed
    crtc_state->adjusted_mode->crtc_clock accordingly, but clk_set_rate()
    used crtc_state->adjusted_mode->clock instead.

    Fixes: 28ab7d35b6e0 ("drm/ingenic: Properly compute timings when using a 3x8-bit panel")
```

```
    Cc: stable@vger.kernel.org # v5.10
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Tested-by: H. Nikolaus Schaller <hns@goldelico.com>     # CI20/jz4780 (HDMI) and
 Alpha400/jz4730 (LCD)
    Acked-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210323144008.166248-1-
 paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 858d38315fc305a3d2a334d1cc7481cb1bb4a367
 Author: Harry Wentland <harry.wentland@amd.com>
 Date:   Thu Apr 22 19:10:52 2021 -0400

    drm/amd/display: Reject non-zero src_y and src_x for video planes

    commit c6c6a712199ab355ce333fa5764a59506bb107c1 upstream.

    [Why]
    This hasn't been well tested and leads to complete system hangs on DCN1
    based systems, possibly others.

    The system hang can be reproduced by gesturing the video on the YouTube
    Android app on ChromeOS into full screen.

    [How]
    Reject atomic commits with non-zero drm_plane_state.src_x or src_y values.

    v2:
     - Add code comment describing the reason we're rejecting non-zero
       src_x and src_y
     - Drop gerrit Change-Id
     - Add stable CC
     - Based on amd-staging-drm-next

    v3: removed trailing whitespace

    Signed-off-by: Harry Wentland <harry.wentland@amd.com>
    Cc: stable@vger.kernel.org
    Cc: nicholas.kazlauskas@amd.com
    Cc: amd-gfx@lists.freedesktop.org
    Cc: alexander.deucher@amd.com
    Cc: Roman.Li@amd.com
    Cc: hersenxs.wu@amd.com
    Cc: danny.wang@amd.com
    Reviewed-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Acked-by: Christian König <christian.koenig@amd.com>
    Reviewed-by: Hersen Wu <hersenxs.wu@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 6683128d7ca548d9138e41381afaca96245805f9
 Author: Maximilian Luz <luzmaximilian@gmail.com>
 Date:   Wed May 12 23:03:16 2021 +0200

    pinctrl/amd: Add device HID for new AMD GPIO controller

    commit 1ca46d3e43569186bd1decfb02a6b4c4ddb4304b upstream.

    Add device HID AMDI0031 to the AMD GPIO controller driver match table.
    This controller can be found on Microsoft Surface Laptop 4 devices and
    seems similar enough that we can just copy the existing AMDI0030 entry.

    Cc: <stable@vger.kernel.org> # 5.10+
    Tested-by: Sachi King <nakato@nakato.io>
    Signed-off-by: Maximilian Luz <luzmaximilian@gmail.com>
```

```
    Link: https://lore.kernel.org/r/20210512210316.1982416-1-luzmaximilian@gmail.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dcb7f78b876c0eedfea68211d6e9472efce95330
Author: Guchun Chen <guchun.chen@amd.com>
Date:   Mon Jun 28 17:03:48 2021 +0800

    drm/amd/display: fix incorrrect valid irq check

    commit e38ca7e422791a4d1c01e56dbf7f9982db0ed365 upstream.

    valid DAL irq should be < DAL_IRQ_SOURCES_NUMBER.

    Signed-off-by: Guchun Chen <guchun.chen@amd.com>
    Reviewed-and-tested-by: Evan Quan <evan.quan@amd.com>
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 72cba7be364868aeb0993f82d9209ce3fe4eadec
Author: Thomas Hebb <tommyhebb@gmail.com>
Date:   Sun Apr 18 19:03:04 2021 -0700

    drm/rockchip: dsi: remove extra component_del() call

    commit b354498bbe65c917d521b3b56317ddc9ab217425 upstream.

    commit cf6d100dd238 ("drm/rockchip: dsi: add dual mipi support") added
    this devcnt field and call to component_del(). However, these both
    appear to be erroneous changes left over from an earlier version of the
    patch. In the version merged, nothing ever modifies devcnt, meaning
    component_del() runs unconditionally and in addition to the
    component_del() calls in dw_mipi_dsi_rockchip_host_detach(). The second
    call fails to delete anything and produces a warning in dmesg.

    If we look at the previous version of the patch[1], however, we see that
    it had logic to calculate devcnt and call component_add() in certain
    situations. This was removed in v6, and the fact that the deletion code
    was not appears to have been an oversight.

    [1] https://patchwork.kernel.org/project/dri-devel/patch/20180821140515.22246-8-
heiko@sntech.de/

    Fixes: cf6d100dd238 ("drm/rockchip: dsi: add dual mipi support")
    Cc: stable@vger.kernel.org
    Signed-off-by: Thomas Hebb <tommyhebb@gmail.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/201385acb0eeb5dfb037afdc6a94bfbcdab97f99.16187
97778.git.tommyhebb@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e7762bc634c8d9b2c1c21e4d39b08a27158aa85b
Author: Lyude Paul <lyude@redhat.com>
Date:   Fri Apr 30 18:34:27 2021 -0400

    drm/dp: Handle zeroed port counts in drm_dp_read_downstream_info()

    commit 205bb69a90363541a634a662a599fddb95956524 upstream.

    While the DP specification isn't entirely clear on if this should be
    allowed or not, some branch devices report having downstream ports present
    while also reporting a downstream port count of 0. So to avoid breaking
```

```
    those devices, we need to handle this in drm_dp_read_downstream_info().

    So, to do this we assume there's no downstream port info when the
    downstream port count is 0.

    Signed-off-by: Lyude Paul <lyude@redhat.com>
    Tested-by: Jérôme de Bretagne <jerome.debretagne@gmail.com>
    Bugzilla: https://gitlab.freedesktop.org/drm/intel/-/issues/3416
    Fixes: 3d3721ccb18a ("drm/i915/dp: Extract drm_dp_read_downstream_info()")
    Cc: <stable@vger.kernel.org> # v5.10+
    Reviewed-by: Ville Syrjälä <ville.syrjala@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210430223428.10514-1-
lyude@redhat.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bab3618ca92cfd19082e50b3b6ac6dfb973865b1
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Fri May 7 17:05:09 2021 +0200

    drm/vc4: hdmi: Prevent clock unbalance

    commit 5b006000423667ef0f55721fc93e477b31f22d28 upstream.

    Since we fixed the hooks to disable the encoder at boot, we now have an
    unbalanced clk_disable call at boot since we never enabled them in the
    first place.

    Let's mimic the state of the hardware and enable the clocks at boot if
    the controller is enabled to get the use-count right.

    Cc: <stable@vger.kernel.org> # v5.10+
    Fixes: 09c438139b8f ("drm/vc4: hdmi: Implement finer-grained hooks")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210507150515.257424-7-
maxime@cerno.tech
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dee39749424bcab86defbd20470b7149ee4584e8
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Fri May 7 17:05:05 2021 +0200

    drm/vc4: crtc: Skip the TXP

    commit 47a50743031ad4138050ae6d266ddd3dfe845ead upstream.

    The vc4_set_crtc_possible_masks is meant to run over all the encoders
    and then set their possible_crtcs mask to their associated pixelvalve.

    However, since the commit 39fcb2808376 ("drm/vc4: txp: Turn the TXP into
    a CRTC of its own"), the TXP has been turned to a CRTC and encoder of
    its own, and while it does indeed register an encoder, it no longer has
    an associated pixelvalve. The code will thus run over the TXP encoder
    and set a bogus possible_crtcs mask, overriding the one set in the TXP
    bind function.

    In order to fix this, let's skip any virtual encoder.

    Cc: <stable@vger.kernel.org> # v5.9+
    Fixes: 39fcb2808376 ("drm/vc4: txp: Turn the TXP into a CRTC of its own")
    Acked-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210507150515.257424-3-
maxime@cerno.tech
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit da442a77dd496e8747ee7457cb70877d66484eba
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Fri May 7 17:05:04 2021 +0200

    drm/vc4: txp: Properly set the possible_crtcs mask

    commit bf6de8e61509f3c957d7f75f017b18d40a18a950 upstream.

    The current code does a binary OR on the possible_crtcs variable of the
    TXP encoder, while we want to set it to that value instead.

    Cc: <stable@vger.kernel.org> # v5.9+
    Fixes: 39fcb2808376 ("drm/vc4: txp: Turn the TXP into a CRTC of its own")
    Acked-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210507150515.257424-2-
maxime@cerno.tech
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e4f2d5c772050a6ec3e08773325798de55438e83
Author: Tiezhu Yang <yangtiezhu@loongson.cn>
Date:   Mon Jun 28 10:53:06 2021 -0400

    drm/radeon: Call radeon_suspend_kms() in radeon_pci_shutdown() for Loongson64

    commit c1bfd74bfef77bcefc88d12eaf8996c0dfd51331 upstream.

    On the Loongson64 platform used with Radeon GPU, shutdown or reboot failed
    when console=tty is in the boot cmdline.

    radeon_suspend_kms() puts the hw in the suspend state, especially set fb
    state as FBINFO_STATE_SUSPENDED:

            if (fbcon) {
                    console_lock();
                    radeon_fbdev_set_suspend(rdev, 1);
                    console_unlock();
            }

    Then avoid to do any more fb operations in the related functions:

            if (p->state != FBINFO_STATE_RUNNING)
                    return;

    So call radeon_suspend_kms() in radeon_pci_shutdown() for Loongson64 to fix
    this issue, it looks like some kind of workaround like powerpc.

    Co-developed-by: Jianmin Lv <lvjianmin@loongson.cn>
    Signed-off-by: Jianmin Lv <lvjianmin@loongson.cn>
    Signed-off-by: Tiezhu Yang <yangtiezhu@loongson.cn>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 63de459218223d59209ddd6a83a9e09490c092ef
Author: Jing Xiangfeng <jingxiangfeng@huawei.com>
Date:   Tue Jun 29 19:44:55 2021 +0800

    drm/radeon: Add the missed drm_gem_object_put() in radeon_user_framebuffer_create()

    commit 9ba85914c36c8fed9bf3e8b69c0782908c1247b7 upstream.

    radeon_user_framebuffer_create() misses to call drm_gem_object_put() in
    an error path. Add the missed function call to fix it.
```

```
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: Jing Xiangfeng <jingxiangfeng@huawei.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7764d2c698d7ab3e38092c808699f05a0e45210d
Author: Evan Quan <evan.quan@amd.com>
Date:   Tue May 25 14:36:29 2021 +0800

    drm/amdgpu: fix the hang caused by PCIe link width switch

    commit adcf949e664a8b04df2fb8aa916892e58561653c upstream.

    SMU had set all the necessary fields for a link width switch
    but the width switch wasn't occurring because the link was idle
    in the L1 state. Setting LC_L1_RECONFIG_EN=0x1 will allow width
    switches to also be initiated while in L1 instead of waiting until
    the link is back in L0.

    Signed-off-by: Evan Quan <evan.quan@amd.com>
    Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0d3a1e59fd9d13d0d0ad586fb7abf3df20ba417c
Author: Evan Quan <evan.quan@amd.com>
Date:   Tue May 25 12:08:53 2021 +0800

    drm/amdgpu: fix NAK-G generation during PCI-e link width switch

    commit 5a5da8ae9546031e43efd4fa5aa8baa481e83dfb upstream.

    A lot of NAK-G being generated when link widht switching is happening.
    WA for this issue is to program the SPC to 4 symbols per clock during
    bootup when the native PCIE width is x4.

    Signed-off-by: Evan Quan <evan.quan@amd.com>
    Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5af51c7bc675637da0e934d1aefda89a5f8503db
Author: Aaron Liu <aaron.liu@amd.com>
Date:   Fri Jun 25 13:50:19 2021 +0800

    drm/amdgpu: enable sdma0 tmz for Raven/Renoir(V2)

    commit e2329e74a615cc58b25c42b7aa1477a5e3f6a435 upstream.

    Without driver loaded, SDMA0_UTCL1_PAGE.TMZ_ENABLE is set to 1
    by default for all asic. On Raven/Renoir, the sdma goldsetting
    changes SDMA0_UTCL1_PAGE.TMZ_ENABLE to 0.
    This patch restores SDMA0_UTCL1_PAGE.TMZ_ENABLE to 1.

    Signed-off-by: Aaron Liu <aaron.liu@amd.com>
    Acked-by: Luben Tuikov <luben.tuikov@amd.com>
    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit fb5beb2b3f06775e84abed6472ea71efb74a4a8c
Author: Joseph Greathouse <Joseph.Greathouse@amd.com>
Date:   Tue Jun 29 21:08:52 2021 -0500

    drm/amdgpu: Update NV SIMD-per-CU to 2

    commit aa6158112645aae514982ad8d56df64428fcf203 upstream.

    Navi series GPUs have 2 SIMDs per CU (and then 2 CUs per WGP).
    The NV enum headers incorrectly listed this as 4, which later meant
    we were incorrectly reporting the number of SIMDs in the HSA
    topology. This could cause problems down the line for user-space
    applications that want to launch a fixed amount of work to each
    SIMD.

    Signed-off-by: Joseph Greathouse <Joseph.Greathouse@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Reviewed-by: Felix Kuehling <Felix.Kuehling@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 788c261cdb58b220b91b3eb61b230fa7af802d19
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Mon Jun 28 15:56:01 2021 -0400

    drm/amdgpu: add new dimgrey cavefish DID

    commit 06ac9b6c736ac9da600b1782d7ac6d6e746286c4 upstream.

    Add new PCI device id.

    Reviewed-by: Guchun Chen <guchun.chen@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7fad97798d9b27efaf945dfa1fcad113f35943b7
Author: Haren Myneni <haren@linux.ibm.com>
Date:   Thu Jun 17 13:29:05 2021 -0700

    powerpc/powernv/vas: Release reference to tgid during window close

    commit 91cdbb955aa94ee0841af4685be40937345d29b8 upstream.

    The kernel handles the NX fault by updating CSB or sending
    signal to process. In multithread applications, children can
    open VAS windows and can exit without closing them. But the
    parent can continue to send NX requests with these windows. To
    prevent pid reuse, reference will be taken on pid and tgid
    when the window is opened and release them during window close.

    The current code is not releasing the tgid reference which can
    cause pid leak and this patch fixes the issue.

    Fixes: db1c08a740635 ("powerpc/vas: Take reference to PID and mm for user space
windows")
    Cc: stable@vger.kernel.org # 5.8+
    Reported-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Haren Myneni <haren@linux.ibm.com>
    Reviewed-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/6020fc4d444864fe20f7dcdc5edfe53e67480a1c.camel@linux.ibm.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 30a401a03189dc06665f97fe6dc0540993abe215
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Fri May 28 11:27:52 2021 -0700

    powerpc/barrier: Avoid collision with clang's __lwsync macro

    commit 015d98149b326e0f1f02e44413112ca8b4330543 upstream.

    A change in clang 13 results in the __lwsync macro being defined as
    __builtin_ppc_lwsync, which emits 'lwsync' or 'msync' depending on what
    the target supports. This breaks the build because of -Werror in
    arch/powerpc, along with thousands of warnings:

     In file included from arch/powerpc/kernel/pmc.c:12:
     In file included from include/linux/bug.h:5:
     In file included from arch/powerpc/include/asm/bug.h:109:
     In file included from include/asm-generic/bug.h:20:
     In file included from include/linux/kernel.h:12:
     In file included from include/linux/bitops.h:32:
     In file included from arch/powerpc/include/asm/bitops.h:62:
     arch/powerpc/include/asm/barrier.h:49:9: error: '__lwsync' macro redefined [-Werror,-
Wmacro-redefined]
     #define __lwsync()      __asm__ __volatile__ (stringify_in_c(LWSYNC) : : :"memory")
             ^
     <built-in>:308:9: note: previous definition is here
     #define __lwsync __builtin_ppc_lwsync
             ^
     1 error generated.

    Undefine this macro so that the runtime patching introduced by
    commit 2d1b2027626d ("powerpc: Fixup lwsync at runtime") continues to
    work properly with clang and the build no longer breaks.

    Cc: stable@vger.kernel.org
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://github.com/ClangBuiltLinux/linux/issues/1386
    Link: https://github.com/llvm/llvm-
project/commit/62b5df7fe2b3fda1772befeda15598fbef96a614
    Link: https://lore.kernel.org/r/20210528182752.1852002-1-nathan@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cc7442037062c8201f34459a3d885ee4078e4d38
Author: Cédric Le Goater <clg@kaod.org>
Date:   Thu Jul 1 17:24:12 2021 +0200

    powerpc/xive: Fix error handling when allocating an IPI

    commit 3f601608b71c3ca1e199898cd16f09d707fedb56 upstream.

    This is a smatch warning:

      arch/powerpc/sysdev/xive/common.c:1161 xive_request_ipi() warn: unsigned 'xid->irq' is
never less than zero.

    Fixes: fd6db2892eba ("powerpc/xive: Modernize XIVE-IPI domain with an 'alloc' handler")
    Cc: stable@vger.kernel.org # v5.13
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Cédric Le Goater <clg@kaod.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20210701152412.1507612-1-clg@kaod.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 7feb683793eadb04ef9b233ee771f1dad77a1781
Author: Naveen N. Rao <naveen.n.rao@linux.vnet.ibm.com>
Date:   Thu Jul 1 20:38:59 2021 +0530

    powerpc/bpf: Reject atomic ops in ppc32 JIT

    commit 307e5042c7bdae15308ef2e9b848833b84122eb0 upstream.

    Commit 91c960b0056672 ("bpf: Rename BPF_XADD and prepare to encode other
    atomics in .imm") converted BPF_XADD to BPF_ATOMIC and updated all JIT
    implementations to reject JIT'ing instructions with an immediate value
    different from BPF_ADD. However, ppc32 BPF JIT was implemented around
    the same time and didn't include the same change. Update the ppc32 JIT
    accordingly.

    Fixes: 51c66ad849a7 ("powerpc/bpf: Implement extended BPF on PPC32")
    Cc: stable@vger.kernel.org # v5.13+
    Signed-off-by: Naveen N. Rao <naveen.n.rao@linux.vnet.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/426699046d89fe50f66ecf74bd31c01eda976ba5.1625145429.git.naveen.n.r
ao@linux.vnet.ibm.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8a96ec5ebf96ad8e2ba7b1b34103a0be5140fc70
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Thu Jul 1 11:17:08 2021 +0000

    powerpc/mm: Fix lockup on kernel exec fault

    commit cd5d5e602f502895e47e18cd46804d6d7014e65c upstream.

    The powerpc kernel is not prepared to handle exec faults from kernel.
    Especially, the function is_exec_fault() will return 'false' when an
    exec fault is taken by kernel, because the check is based on reading
    current->thread.regs->trap which contains the trap from user.

    For instance, when provoking a LKDTM EXEC_USERSPACE test,
    current->thread.regs->trap is set to SYSCALL trap (0xc00), and
    the fault taken by the kernel is not seen as an exec fault by
    set_access_flags_filter().

    Commit d7df2443cd5f ("powerpc/mm: Fix spurious segfaults on radix
    with autonuma") made it clear and handled it properly. But later on
    commit d3ca587404b3 ("powerpc/mm: Fix reporting of kernel execute
    faults") removed that handling, introducing test based on error_code.
    And here is the problem, because on the 603 all upper bits of SRR1
    get cleared when the TLB instruction miss handler bails out to ISI.

    Until commit cbd7e6ca0210 ("powerpc/fault: Avoid heavy
    search_exception_tables() verification"), an exec fault from kernel
    at a userspace address was indirectly caught by the lack of entry for
    that address in the exception tables. But after that commit the
    kernel mainly relies on KUAP or on core mm handling to catch wrong
    user accesses. Here the access is not wrong, so mm handles it.
    It is a minor fault because PAGE_EXEC is not set,
    set_access_flags_filter() should set PAGE_EXEC and voila.
    But as is_exec_fault() returns false as explained in the beginning,
    set_access_flags_filter() bails out without setting PAGE_EXEC flag,
    which leads to a forever minor exec fault.

    As the kernel is not prepared to handle such exec faults, the thing to
    do is to fire in bad_kernel_fault() for any exec fault taken by the
    kernel, as it was prior to commit d3ca587404b3.
```

```
    Fixes: d3ca587404b3 ("powerpc/mm: Fix reporting of kernel execute faults")
    Cc: stable@vger.kernel.org # v4.14+
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Acked-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/024bb05105050f704743a0083fe3548702be5706.1625138205.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 987a852734e668e0829cf65dd182e44a03064800
Author: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
Date:   Wed Jul 7 18:10:15 2021 -0700

    mm/mremap: hold the rmap lock in write mode when moving page table entries.

    commit 97113eb39fa7972722ff490b947d8af023e1f6a2 upstream.

    To avoid a race between rmap walk and mremap, mremap does
    take_rmap_locks().  The lock was taken to ensure that rmap walk don't miss
    a page table entry due to PTE moves via move_pagetables().  The kernel
    does further optimization of this lock such that if we are going to find
    the newly added vma after the old vma, the rmap lock is not taken.  This
    is because rmap walk would find the vmas in the same order and if we don't
    find the page table attached to older vma we would find it with the new
    vma which we would iterate later.

    As explained in commit eb66ae030829 ("mremap: properly flush TLB before
    releasing the page") mremap is special in that it doesn't take ownership
    of the page.  The optimized version for PUD/PMD aligned mremap also
    doesn't hold the ptl lock.  This can result in stale TLB entries as show
    below.

    This patch updates the rmap locking requirement in mremap to handle the race condition
    explained below with optimized mremap::

    Optmized PMD move

        CPU 1                           CPU 2                                   CPU 3

        mremap(old_addr, new_addr)      page_shrinker/try_to_unmap_one

        mmap_write_lock_killable()

                                        addr = old_addr
                                        lock(pte_ptl)
        lock(pmd_ptl)
        pmd = *old_pmd
        pmd_clear(old_pmd)
        flush_tlb_range(old_addr)

        *new_pmd = pmd
                                                                                *new_addr =
10; and fills
                                                                                TLB with new
addr
                                                                                and old pfn

        unlock(pmd_ptl)
                                        ptep_clear_flush()
                                        old pfn is free.
                                                                                Stale TLB
entry

    Optimized PUD move also suffers from a similar race.  Both the above race
```

```
    condition can be fixed if we force mremap path to take rmap lock.

    Link: https://lkml.kernel.org/r/20210616045239.370802-7-aneesh.kumar@linux.ibm.com
    Fixes: 2c91bd4a4e2e ("mm: speed up mremap by 20x on large regions")
    Fixes: c49dd3401802 ("mm: speedup mremap on 1GB or larger regions")
    Link: https://lore.kernel.org/linux-mm/CAHk-=wgXVR04eBNtxQfevontWnP6FDm+oj5vauQXP3S-
huwbPw@mail.gmail.com
    Signed-off-by: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
    Acked-by: Hugh Dickins <hughd@google.com>
    Acked-by: Kirill A. Shutemov <kirill.shutemov@linux.intel.com>
    Cc: Christophe Leroy <christophe.leroy@csgroup.eu>
    Cc: Joel Fernandes <joel@joelfernandes.org>
    Cc: Kalesh Singh <kaleshsingh@google.com>
    Cc: Kirill A. Shutemov <kirill@shutemov.name>
    Cc: Michael Ellerman <mpe@ellerman.id.au>
    Cc: Nicholas Piggin <npiggin@gmail.com>
    Cc: Stephen Rothwell <sfr@canb.auug.org.au>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f908381a3ec721fb259bfbc6cbe4fdb607e074c2
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Fri Jun 25 11:49:29 2021 +0100

    MIPS: MT extensions are not available on MIPS32r1

    commit cad065ed8d8831df67b9754cc4437ed55d8b48c0 upstream.

    MIPS MT extensions were added with the MIPS 34K processor, which was
    based on the MIPS32r2 ISA.

    This fixes a build error when building a generic kernel for a MIPS32r1
    CPU.

    Fixes: c434b9f80b09 ("MIPS: Kconfig: add MIPS_GENERIC_KERNEL symbol")
    Cc: stable@vger.kernel.org # v5.9
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eecc96e0458c54809618277f6565b72581541bd4
Author: Vidya Sagar <vidyas@nvidia.com>
Date:   Tue May 4 22:51:57 2021 +0530

    PCI: tegra194: Fix host initialization during resume

    commit c4bf1f25c6c187864681d5ad4dd1fa92f62d5d32 upstream.

    Commit 275e88b06a27 ("PCI: tegra: Fix host link initialization") broke
    host initialization during resume as it misses out calling the API
    dw_pcie_setup_rc() which is required for host and MSI initialization.

    Link: https://lore.kernel.org/r/20210504172157.29712-1-vidyas@nvidia.com
    Fixes: 275e88b06a27 ("PCI: tegra: Fix host link initialization")
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Vidya Sagar <vidyas@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b63416bc9d6de0d9a8d0c66c3e816d5a4d5f5d98
Author: 周琰杰 (Zhou Yanjie) <zhouyanjie@wanyeetech.com>
Date:   Sat Jun 26 14:18:40 2021 +0800
```

```
    MIPS: CI20: Reduce clocksource to 750 kHz.

    [ Upstream commit 23c64447b3538a6f34cb38aae3bc19dc1ec53436 ]

    The original clock (3 MHz) is too fast for the clocksource,
    there will be a chance that the system may get stuck.

    Reported-by: Nikolaus Schaller <hns@goldelico.com>
    Tested-by: Nikolaus Schaller <hns@goldelico.com> # on CI20
    Signed-off-by: 周琰杰 (Zhou Yanjie) <zhouyanjie@wanyeetech.com>
    Acked-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e3f6e978e4c69d23693caf4c84beb6ee9a297ba
Author: Nick Desaulniers <ndesaulniers@google.com>
Date:   Mon Jun 28 14:50:26 2021 -0700

    MIPS: set mips32r5 for virt extensions

    [ Upstream commit c994a3ec7ecc8bd2a837b2061e8a76eb8efc082b ]

    Clang's integrated assembler only accepts these instructions when the
    cpu is set to mips32r5. With this change, we can assemble
    malta_defconfig with Clang via `make LLVM_IAS=1`.

    Link: https://github.com/ClangBuiltLinux/linux/issues/763
    Reported-by: Dmitry Golovin <dima@golovin.in>
    Signed-off-by: Nick Desaulniers <ndesaulniers@google.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1930d84a29d41a33c78a29ec0bec3e62cd57f06c
Author: zhanglianjie <zhanglianjie@uniontech.com>
Date:   Mon Jun 28 17:11:05 2021 +0800

    MIPS: loongsoon64: Reserve memory below starting pfn to prevent Oops

    [ Upstream commit 6817c944430d00f71ccaa9c99ff5b0096aeb7873 ]

    The cause of the problem is as follows:
    1. when cat /sys/devices/system/memory/memory0/valid_zones,
       test_pages_in_a_zone() will be called.
    2. test_pages_in_a_zone() finds the zone according to stat_pfn = 0.
       The smallest pfn of the numa node in the mips architecture is 128,
       and the page corresponding to the previous 0~127 pfn is not
       initialized (page->flags is 0xFFFFFFFF)
    3. The nid and zonenum obtained using page_zone(pfn_to_page(0)) are out
       of bounds in the corresponding array,
       &NODE_DATA(page_to_nid(page))->node_zones[page_zonenum(page)],
       access to the out-of-bounds zone member variables appear abnormal,
       resulting in Oops.
    Therefore, it is necessary to keep the page between 0 and the minimum
    pfn to prevent Oops from appearing.

    Signed-off-by: zhanglianjie <zhanglianjie@uniontech.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92b74375a1bbf5f7f64f4ea98064a5ba62956bcc
Author: Marcelo Ricardo Leitner <marcelo.leitner@gmail.com>
Date:   Mon Jun 28 16:13:42 2021 -0300

    sctp: add size validation when walking chunks
```

```
    [ Upstream commit 50619dbf8db77e98d821d615af4f634d08e22698 ]

    The first chunk in a packet is ensured to be present at the beginning of
    sctp_rcv(), as a packet needs to have at least 1 chunk. But the second
    one, may not be completely available and ch->length can be over
    uninitialized memory.

    Fix here is by only trying to walk on the next chunk if there is enough to
    hold at least the header, and then proceed with the ch->length validation
    that is already there.

    Reported-by: Ilja Van Sprundel <ivansprundel@ioactive.com>
    Signed-off-by: Marcelo Ricardo Leitner <marcelo.leitner@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ecabee69d190f2bd9bdc5140109a27231428413
Author: Marcelo Ricardo Leitner <marcelo.leitner@gmail.com>
Date:   Mon Jun 28 16:13:41 2021 -0300

    sctp: validate from_addr_param return

    [ Upstream commit 0c5dc070ff3d6246d22ddd931f23a6266249e3db ]

    Ilja reported that, simply putting it, nothing was validating that
    from_addr_param functions were operating on initialized memory. That is,
    the parameter itself was being validated by sctp_walk_params, but it
    doesn't check for types and their specific sizes and it could be a 0-length
    one, causing from_addr_param to potentially work over the next parameter or
    even uninitialized memory.

    The fix here is to, in all calls to from_addr_param, check if enough space
    is there for the wanted IP address type.

    Reported-by: Ilja Van Sprundel <ivansprundel@ioactive.com>
    Signed-off-by: Marcelo Ricardo Leitner <marcelo.leitner@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f7ca3459118382d2d061dcc09b184ac6ce20ef6
Author: gushengxian <gushengxian@yulong.com>
Date:   Sat Jun 26 04:56:06 2021 -0700

    flow_offload: action should not be NULL when it is referenced

    [ Upstream commit 9ea3e52c5bc8bb4a084938dc1e3160643438927a ]

    "action" should not be NULL when it is referenced.

    Signed-off-by: gushengxian <13145886936@163.com>
    Signed-off-by: gushengxian <gushengxian@yulong.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9d02dbbab0ad73f5727ebbf050aac27f00ce105
Author: Rustam Kovhaev <rkovhaev@gmail.com>
Date:   Sat Jun 26 11:11:56 2021 -0700

    bpf: Fix false positive kmemleak report in bpf_ringbuf_area_alloc()

    [ Upstream commit ccff81e1d028bbbf8573d3364a87542386c707bf ]

    kmemleak scans struct page, but it does not scan the page content. If we
    allocate some memory with kmalloc(), then allocate page with alloc_page(),
```

```
    and if we put kmalloc pointer somewhere inside that page, kmemleak will
    report kmalloc pointer as a false positive.

    We can instruct kmemleak to scan the memory area by calling kmemleak_alloc()
    and kmemleak_free(), but part of struct bpf_ringbuf is mmaped to user space,
    and if struct bpf_ringbuf changes we would have to revisit and review size
    argument in kmemleak_alloc(), because we do not want kmemleak to scan the
    user space memory. Let's simplify things and use kmemleak_not_leak() here.

    For posterity, also adding additional prior analysis from Andrii:

      I think either kmemleak or syzbot are misreporting this. I've added a
      bunch of printks around all allocations performed by BPF ringbuf. [...]
      On repro side I get these two warnings:

      [vmuser@archvm bpf]$ sudo ./repro
      BUG: memory leak
      unreferenced object 0xffff88810d538c00 (size 64):
        comm "repro", pid 2140, jiffies 4294692933 (age 14.540s)
        hex dump (first 32 bytes):
          00 af 19 04 00 ea ff ff c0 ae 19 04 00 ea ff ff  ................
          80 ae 19 04 00 ea ff ff c0 29 2e 04 00 ea ff ff  .........)......
        backtrace:
          [<0000000077bfbfbd>] __bpf_map_area_alloc+0x31/0xc0
          [<00000000587fa522>] ringbuf_map_alloc.cold.4+0x48/0x218
          [<0000000044d49e96>] __do_sys_bpf+0x359/0x1d90
          [<00000000f601d565>] do_syscall_64+0x2d/0x40
          [<0000000043d3112a>] entry_SYSCALL_64_after_hwframe+0x44/0xae

      BUG: memory leak
      unreferenced object 0xffff88810d538c80 (size 64):
        comm "repro", pid 2143, jiffies 4294699025 (age 8.448s)
        hex dump (first 32 bytes):
          80 aa 19 04 00 ea ff ff 00 ab 19 04 00 ea ff ff  ................
          c0 ab 19 04 00 ea ff ff 80 44 28 04 00 ea ff ff  .........D(.....
        backtrace:
          [<0000000077bfbfbd>] __bpf_map_area_alloc+0x31/0xc0
          [<00000000587fa522>] ringbuf_map_alloc.cold.4+0x48/0x218
          [<0000000044d49e96>] __do_sys_bpf+0x359/0x1d90
          [<00000000f601d565>] do_syscall_64+0x2d/0x40
          [<0000000043d3112a>] entry_SYSCALL_64_after_hwframe+0x44/0xae

      Note that both reported leaks (ffff88810d538c80 and ffff88810d538c00)
      correspond to pages array bpf_ringbuf is allocating and tracking properly
      internally. Note also that syzbot repro doesn't close FD of created BPF
      ringbufs, and even when ./repro itself exits with error, there are still
      two forked processes hanging around in my system. So clearly ringbuf maps
      are alive at that point. So reporting any memory leak looks weird at that
      point, because that memory is being used by active referenced BPF ringbuf.

      It's also a question why repro doesn't clean up its forks. But if I do a
      `pkill repro`, I do see that all the allocated memory is /properly/ cleaned
      up [and the] "leaks" are deallocated properly.

      BTW, if I add close() right after bpf() syscall in syzbot repro, I see that
      everything is immediately deallocated, like designed. And no memory leak
      is reported. So I don't think the problem is anywhere in bpf_ringbuf code,
      rather in the leak detection and/or repro itself.

    Reported-by: syzbot+5d895828587f49e7fe9b@syzkaller.appspotmail.com
    Signed-off-by: Rustam Kovhaev <rkovhaev@gmail.com>
    [ Daniel: also included analysis from Andrii to the commit log ]
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Tested-by: syzbot+5d895828587f49e7fe9b@syzkaller.appspotmail.com
    Cc: Dmitry Vyukov <dvyukov@google.com>
```

```
    Cc: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/CAEf4BzYk+dqs+jwu6VKXP-
RttcTEGFe+ySTGWT9CRNkagDiJVA@mail.gmail.com
    Link: https://lore.kernel.org/lkml/YNTAqiE7CWJhOK2M@nuc10
    Link: https://lore.kernel.org/lkml/20210615101515.GC26027@arm.com
    Link: https://syzkaller.appspot.com/bug?extid=5d895828587f49e7fe9b
    Link: https://lore.kernel.org/bpf/20210626181156.1873604-1-rkovhaev@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f2d7f0241e220520318d79d14151fbafa5194db
Author: Odin Ugedal <odin@uged.al>
Date:   Thu Jun 24 13:18:15 2021 +0200

    sched/fair: Ensure _sum and _avg values stay consistent

    [ Upstream commit 1c35b07e6d3986474e5635be566e7bc79d97c64d ]

    The _sum and _avg values are in general sync together with the PELT
    divider. They are however not always completely in perfect sync,
    resulting in situations where _sum gets to zero while _avg stays
    positive. Such situations are undesirable.

    This comes from the fact that PELT will increase period_contrib, also
    increasing the PELT divider, without updating _sum and _avg values to
    stay in perfect sync where (_sum == _avg * divider). However, such PELT
    change will never lower _sum, making it impossible to end up in a
    situation where _sum is zero and _avg is not.

    Therefore, we need to ensure that when subtracting load outside PELT,
    that when _sum is zero, _avg is also set to zero. This occurs when
    (_sum < _avg * divider), and the subtracted (_avg * divider) is bigger
    or equal to the current _sum, while the subtracted _avg is smaller than
    the current _avg.

    Reported-by: Sachin Sant <sachinp@linux.vnet.ibm.com>
    Reported-by: Naresh Kamboju <naresh.kamboju@linaro.org>
    Signed-off-by: Odin Ugedal <odin@uged.al>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Vincent Guittot <vincent.guittot@linaro.org>
    Tested-by: Sachin Sant <sachinp@linux.vnet.ibm.com>
    Link: https://lore.kernel.org/r/20210624111815.57937-1-odin@uged.al
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 377c584c325efe4eca9d7117684a07e82267213d
Author: Tim Jiang <tjiang@codeaurora.org>
Date:   Tue Jun 1 17:57:10 2021 +0800

    Bluetooth: btusb: fix bt fiwmare downloading failure issue for qca btsoc.

    [ Upstream commit 4f00bfb372674d586c4a261bfc595cbce101fbb6 ]

    This is btsoc timing issue, after host start to downloading bt firmware,
    ep2 need time to switch from function acl to function dfu, so host add
    20ms delay as workaround.

    Signed-off-by: Tim Jiang <tjiang@codeaurora.org>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 387d99b463afd0f6d6db006b3008eee5c5d68c85
Author: Tim Jiang <tjiang@codeaurora.org>
Date:   Tue Jun 1 17:57:25 2021 +0800

    Bluetooth: btusb: use default nvm if boardID is 0 for wcn6855.
```

```
    [ Upstream commit ca17a5cccf8b6d35dab4729bea8f4350bc0b4caf ]

    if boardID is 0, will use the default nvm file without surfix.

    Signed-off-by: Tim Jiang <tjiang@codeaurora.org>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e4a714f451046223331f65b64c718fa393303ab
Author: Tedd Ho-Jeong An <tedd.an@intel.com>
Date:   Wed May 26 10:36:22 2021 -0700

    Bluetooth: mgmt: Fix the command returns garbage parameter value

    [ Upstream commit 02ce2c2c24024aade65a8d91d6a596651eaf2d0a ]

    When the Get Device Flags command fails, it returns the error status
    with the parameters filled with the garbage values. Although the
    parameters are not used, it is better to fill with zero than the random
    values.

    Signed-off-by: Tedd Ho-Jeong An <tedd.an@intel.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15407b14e27b93b7589e7b6d22e80a3ef6e39085
Author: Hilda Wu <hildawu@realtek.com>
Date:   Fri May 14 11:19:01 2021 +0800

    Bluetooth: btusb: Add support USB ALT 3 for WBS

    [ Upstream commit e848dbd364aca44c9d23c04bef964fab79e2b34f ]

    Because mSBC frames do not need to be aligned to the SCO packet
    boundary. Using USB ALT 3 let HCI payload >= 60 bytes, let mSBC
    data satisfy 60 Bytes avoid payload unaligned situation and fixed
    some headset no voise issue.

    USB Alt 3 supported also need HFP support transparent MTU in 72 Bytes.

    Signed-off-by: Hilda Wu <hildawu@realtek.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69ccde876ca58ad7cd4eb7b4de9311a2d645922d
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed May 19 13:41:51 2021 -0700

    Bluetooth: L2CAP: Fix invalid access on ECRED Connection response

    [ Upstream commit de895b43932cb47e69480540be7eca289af24f23 ]

    The use of l2cap_chan_del is not safe under a loop using
    list_for_each_entry.

    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 840907eecac0e7b5def12130b99aeb0f800c29f6
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed May 19 13:41:50 2021 -0700

    Bluetooth: L2CAP: Fix invalid access if ECRED Reconfigure fails
```

```
    [ Upstream commit 1fa20d7d4aad02206e84b74915819fbe9f81dab3 ]

    The use of l2cap_chan_del is not safe under a loop using
    list_for_each_entry.

    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit efffb7884033b87b8c7ee1e11f7cd73e1c2ffce9
Author: Daniel Lenski <dlenski@gmail.com>
Date:   Tue May 18 10:25:46 2021 -0700

    Bluetooth: btusb: Add a new QCA_ROME device (0cf3:e500)

    [ Upstream commit 0324d19cb99804d99e42c990b8b1e191575a091b ]

    This patch adds the 0cf3:e500 Bluetooth device (from a QCA9377 board) as a
    QCA_ROME device.  It appears to be functionally identical to another device
    ID, also from a QCA9377 board, which was previously marked as QCA_ROME in
    0a03f98b98c201191e3ba15a0e33f46d8660e1fd
    ("Bluetooth: Add a new 04ca:3015 QCA_ROME device").

    Without this patch, the WiFi side of the QCA9377 board is slow or unusable
    when the Bluetooth side is in use.

    See https://askubuntu.com/a/1137852 for another report of QCA_ROME fixing
    this issue for this device ID.

    /sys/kernel/debug/usb/devices:

    T:  Bus=05 Lev=01 Prnt=01 Port=00 Cnt=01 Dev#=  2 Spd=12   MxCh= 0
    D:  Ver= 2.01 Cls=e0(wlcon) Sub=01 Prot=01 MxPS=64 #Cfgs=  1
    P:  Vendor=0cf3 ProdID=e500 Rev= 0.01
    C:* #Ifs= 2 Cfg#= 1 Atr=e0 MxPwr=100mA
    I:* If#= 0 Alt= 0 #EPs= 3 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=81(I) Atr=03(Int.) MxPS=  16 Ivl=1ms
    E:  Ad=82(I) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    E:  Ad=02(O) Atr=02(Bulk) MxPS=  64 Ivl=0ms
    I:* If#= 1 Alt= 0 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=   0 Ivl=1ms
    I:  If#= 1 Alt= 1 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=   9 Ivl=1ms
    I:  If#= 1 Alt= 2 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  17 Ivl=1ms
    I:  If#= 1 Alt= 3 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  25 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  25 Ivl=1ms
    I:  If#= 1 Alt= 4 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  33 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  33 Ivl=1ms
    I:  If#= 1 Alt= 5 #EPs= 2 Cls=e0(wlcon) Sub=01 Prot=01 Driver=btusb
    E:  Ad=83(I) Atr=01(Isoc) MxPS=  49 Ivl=1ms
    E:  Ad=03(O) Atr=01(Isoc) MxPS=  49 Ivl=1ms

    Signed-off-by: Daniel Lenski <dlenski@gmail.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c19a2820df32b94d1d4fb65e96ae36f97dddda56
Author: Kai-Heng Feng <kai.heng.feng@canonical.com>
```

```
Date:   Fri May 14 15:14:52 2021 +0800

    Bluetooth: Shutdown controller after workqueues are flushed or cancelled

    [ Upstream commit 0ea9fd001a14ebc294f112b0361a4e601551d508 ]

    Rfkill block and unblock Intel USB Bluetooth [8087:0026] may make it
    stops working:
    [  509.691509] Bluetooth: hci0: HCI reset during shutdown failed
    [  514.897584] Bluetooth: hci0: MSFT filter_enable is already on
    [  530.044751] usb 3-10: reset full-speed USB device number 5 using xhci_hcd
    [  545.660350] usb 3-10: device descriptor read/64, error -110
    [  561.283530] usb 3-10: device descriptor read/64, error -110
    [  561.519682] usb 3-10: reset full-speed USB device number 5 using xhci_hcd
    [  566.686650] Bluetooth: hci0: unexpected event for opcode 0x0500
    [  568.752452] Bluetooth: hci0: urb 0000000096cd309b failed to resubmit (113)
    [  578.797955] Bluetooth: hci0: Failed to read MSFT supported features (-110)
    [  586.286565] Bluetooth: hci0: urb 00000000c522f633 failed to resubmit (113)
    [  596.215302] Bluetooth: hci0: Failed to read MSFT supported features (-110)

    Or kernel panics because other workqueues already freed skb:
    [ 2048.663763] BUG: kernel NULL pointer dereference, address: 0000000000000000
    [ 2048.663775] #PF: supervisor read access in kernel mode
    [ 2048.663779] #PF: error_code(0x0000) - not-present page
    [ 2048.663782] PGD 0 P4D 0
    [ 2048.663787] Oops: 0000 [#1] SMP NOPTI
    [ 2048.663793] CPU: 3 PID: 4491 Comm: rfkill Tainted: G        W         5.13.0-rc1-
next-20210510+ #20
    [ 2048.663799] Hardware name: HP HP EliteBook 850 G8 Notebook PC/8846, BIOS T76 Ver.
01.01.04 12/02/2020
    [ 2048.663801] RIP: 0010:__skb_ext_put+0x6/0x50
    [ 2048.663814] Code: 8b 1b 48 85 db 75 db 5b 41 5c 5d c3 be 01 00 00 00 e8 de 13 c0 ff
eb e7 be 02 00 00 00 e8 d2 13 c0 ff eb db 0f 1f 44 00 00 55 <8b> 07 48 89 e5 83 f8 01 74 14
b8 ff ff ff ff f0 0f c1
    07 83 f8 01
    [ 2048.663819] RSP: 0018:ffffc1d105b6fd80 EFLAGS: 00010286
    [ 2048.663824] RAX: 0000000000000000 RBX: ffff9d9ac5649000 RCX: 0000000000000000
    [ 2048.663827] RDX: ffffffffc0d1daf6 RSI: 0000000000000206 RDI: 0000000000000000
    [ 2048.663830] RBP: ffffc1d105b6fd98 R08: 0000000000000001 R09: ffff9d9ace8ceac0
    [ 2048.663834] R10: ffff9d9ace8ceac0 R11: 0000000000000001 R12: ffff9d9ac5649000
    [ 2048.663838] R13: 0000000000000000 R14: 00007ffe0354d650 R15: 0000000000000000
    [ 2048.663843] FS:  00007fe02ab19740(0000) GS:ffff9d9e5f8c0000(0000)
knlGS:0000000000000000
    [ 2048.663849] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [ 2048.663853] CR2: 0000000000000000 CR3: 0000000111a52004 CR4: 0000000000770ee0
    [ 2048.663856] PKRU: 55555554
    [ 2048.663859] Call Trace:
    [ 2048.663865]  ? skb_release_head_state+0x5e/0x80
    [ 2048.663873]  kfree_skb+0x2f/0xb0
    [ 2048.663881]  btusb_shutdown_intel_new+0x36/0x60 [btusb]
    [ 2048.663905]  hci_dev_do_close+0x48c/0x5e0 [bluetooth]
    [ 2048.663954]  ? __cond_resched+0x1a/0x50
    [ 2048.663962]  hci_rfkill_set_block+0x56/0xa0 [bluetooth]
    [ 2048.664007]  rfkill_set_block+0x98/0x170
    [ 2048.664016]  rfkill_fop_write+0x136/0x1e0
    [ 2048.664022]  vfs_write+0xc7/0x260
    [ 2048.664030]  ksys_write+0xb1/0xe0
    [ 2048.664035]  ? exit_to_user_mode_prepare+0x37/0x1c0
    [ 2048.664042]  __x64_sys_write+0x1a/0x20
    [ 2048.664048]  do_syscall_64+0x40/0xb0
    [ 2048.664055]  entry_SYSCALL_64_after_hwframe+0x44/0xae
    [ 2048.664060] RIP: 0033:0x7fe02ac23c27
    [ 2048.664066] Code: 0d 00 f7 d8 64 89 02 48 c7 c0 ff ff ff ff eb b7 0f 1f 00 f3 0f 1e
fa 64 8b 04 25 18 00 00 00 85 c0 75 10 b8 01 00 00 00 0f 05 <48> 3d 00 f0 ff ff 77 51 c3 48
83 ec 28 48 89 54 24 18 48 89 74 24
```

```
     [ 2048.664070] RSP: 002b:00007ffe0354d638 EFLAGS: 00000246 ORIG_RAX: 0000000000000001
     [ 2048.664075] RAX: ffffffffffffffda RBX: 0000000000000001 RCX: 00007fe02ac23c27
     [ 2048.664078] RDX: 0000000000000008 RSI: 00007ffe0354d650 RDI: 0000000000000003
     [ 2048.664081] RBP: 0000000000000000 R08: 0000559b05998440 R09: 0000559b05998440
     [ 2048.664084] R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000003
     [ 2048.664086] R13: 0000000000000000 R14: ffffffff00000000 R15: 00000000ffffffff

     So move the shutdown callback to a place where workqueues are either
     flushed or cancelled to resolve the issue.

     Signed-off-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
     Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 286bba24e6ce5db00cc7e5f7dce39ec15033667f
  Author: Kiran K <kiran.k@intel.com>
  Date:   Thu Apr 8 22:31:59 2021 +0530

     Bluetooth: Fix alt settings for incoming SCO with transparent coding format

     [ Upstream commit 06d213d8a89a6f55b708422c3dda2b22add10748 ]

     For incoming SCO connection with transparent coding format, alt setting
     of CVSD is getting applied instead of Transparent.

     Before fix:
     < HCI Command: Accept Synchron.. (0x01|0x0029) plen 21  #2196 [hci0] 321.342548
             Address: 1C:CC:D6:E2:EA:80 (Xiaomi Communications Co Ltd)
             Transmit bandwidth: 8000
             Receive bandwidth: 8000
             Max latency: 13
             Setting: 0x0003
               Input Coding: Linear
               Input Data Format: 1's complement
               Input Sample Size: 8-bit
               # of bits padding at MSB: 0
               Air Coding Format: Transparent Data
             Retransmission effort: Optimize for link quality (0x02)
             Packet type: 0x003f
               HV1 may be used
               HV2 may be used
               HV3 may be used
               EV3 may be used
               EV4 may be used
               EV5 may be used
     > HCI Event: Command Status (0x0f) plen 4               #2197 [hci0] 321.343585
           Accept Synchronous Connection Request (0x01|0x0029) ncmd 1
             Status: Success (0x00)
     > HCI Event: Synchronous Connect Comp.. (0x2c) plen 17  #2198 [hci0] 321.351666
             Status: Success (0x00)
             Handle: 257
             Address: 1C:CC:D6:E2:EA:80 (Xiaomi Communications Co Ltd)
             Link type: eSCO (0x02)
             Transmission interval: 0x0c
             Retransmission window: 0x04
             RX packet length: 60
             TX packet length: 60
             Air mode: Transparent (0x03)
     ........
     > SCO Data RX: Handle 257 flags 0x00 dlen 48            #2336 [hci0] 321.383655
     < SCO Data TX: Handle 257 flags 0x00 dlen 60            #2337 [hci0] 321.389558
     > SCO Data RX: Handle 257 flags 0x00 dlen 48            #2338 [hci0] 321.393615
     > SCO Data RX: Handle 257 flags 0x00 dlen 48            #2339 [hci0] 321.393618
     > SCO Data RX: Handle 257 flags 0x00 dlen 48            #2340 [hci0] 321.393618
     < SCO Data TX: Handle 257 flags 0x00 dlen 60            #2341 [hci0] 321.397070
```

```
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2342 [hci0] 321.403622
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2343 [hci0] 321.403625
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2344 [hci0] 321.403625
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2345 [hci0] 321.403625
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #2346 [hci0] 321.404569
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #2347 [hci0] 321.412091
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2348 [hci0] 321.413626
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2349 [hci0] 321.413630
> SCO Data RX: Handle 257 flags 0x00 dlen 48           #2350 [hci0] 321.413630
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #2351 [hci0] 321.419674

After fix:

< HCI Command: Accept Synchronou.. (0x01|0x0029) plen 21  #309 [hci0] 49.439693
        Address: 1C:CC:D6:E2:EA:80 (Xiaomi Communications Co Ltd)
        Transmit bandwidth: 8000
        Receive bandwidth: 8000
        Max latency: 13
        Setting: 0x0003
          Input Coding: Linear
          Input Data Format: 1's complement
          Input Sample Size: 8-bit
          # of bits padding at MSB: 0
          Air Coding Format: Transparent Data
        Retransmission effort: Optimize for link quality (0x02)
        Packet type: 0x003f
          HV1 may be used
          HV2 may be used
          HV3 may be used
          EV3 may be used
          EV4 may be used
          EV5 may be used
> HCI Event: Command Status (0x0f) plen 4                #310 [hci0] 49.440308
      Accept Synchronous Connection Request (0x01|0x0029) ncmd 1
        Status: Success (0x00)
> HCI Event: Synchronous Connect Complete (0x2c) plen 17  #311 [hci0] 49.449308
        Status: Success (0x00)
        Handle: 257
        Address: 1C:CC:D6:E2:EA:80 (Xiaomi Communications Co Ltd)
        Link type: eSCO (0x02)
        Transmission interval: 0x0c
        Retransmission window: 0x04
        RX packet length: 60
        TX packet length: 60
        Air mode: Transparent (0x03)
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #312 [hci0] 49.450421
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #313 [hci0] 49.457927
> HCI Event: Max Slots Change (0x1b) plen 3            #314 [hci0] 49.460345
        Handle: 256
        Max slots: 5
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #315 [hci0] 49.465453
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #316 [hci0] 49.470502
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #317 [hci0] 49.470519
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #318 [hci0] 49.472996
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #319 [hci0] 49.480412
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #320 [hci0] 49.480492
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #321 [hci0] 49.487989
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #322 [hci0] 49.490303
< SCO Data TX: Handle 257 flags 0x00 dlen 60           #323 [hci0] 49.495496
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #324 [hci0] 49.500304
> SCO Data RX: Handle 257 flags 0x00 dlen 60           #325 [hci0] 49.500311

Signed-off-by: Kiran K <kiran.k@intel.com>
Signed-off-by: Lokendra Singh <lokendra.singh@intel.com>
Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
```

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

commit cbb2f564faa4a351740c1f27d0a45dbe0ab5d24c
Author: Yu Liu <yudiliu@google.com>
Date:   Mon Apr 19 16:53:30 2021 -0700

    Bluetooth: Fix the HCI to MGMT status conversion table

    [ Upstream commit 4ef36a52b0e47c80bbfd69c0cce61c7ae9f541ed ]

    0x2B, 0x31 and 0x33 are reserved for future use but were not present in
    the HCI to MGMT conversion table, this caused the conversion to be
    incorrect for the HCI status code greater than 0x2A.

    Reviewed-by: Miao-chen Chou <mcchou@chromium.org>
    Signed-off-by: Yu Liu <yudiliu@google.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c73a8008035879a27f4bcf2d87869d039cb98fb
Author: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
Date:   Tue Apr 13 13:21:03 2021 -0300

    Bluetooth: cmtp: fix file refcount when cmtp_attach_device fails

    [ Upstream commit 3cfdf8fcaafa62a4123f92eb0f4a72650da3a479 ]

    When cmtp_attach_device fails, cmtp_add_connection returns the error value
    which leads to the caller to doing fput through sockfd_put. But
    cmtp_session kthread, which is stopped in this path will also call fput,
    leading to a potential refcount underflow or a use-after-free.

    Add a refcount before we signal the kthread to stop. The kthread will try
    to grab the cmtp_session_sem mutex before doing the fput, which is held
    when get_file is called, so there should be no races there.

    Reported-by: Ryota Shiga
    Signed-off-by: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc239f449b031e310f2b6aecb6f39e91ed60c67b
Author: mark-yw.chen <mark-yw.chen@mediatek.com>
Date:   Mon Apr 12 23:06:26 2021 +0800

    Bluetooth: btusb: Fixed too many in-token issue for Mediatek Chip.

    [ Upstream commit 8454ed9ff9647e31e061fb5eb2e39ce79bc5e960 ]

    This patch reduce in-token during download patch procedure.
    Don't submit urb for polling event before sending hci command.

    Signed-off-by: mark-yw.chen <mark-yw.chen@mediatek.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07583ba2e2d8947c3d365d97608cb436510885ac
Author: Gerd Rausch <gerd.rausch@oracle.com>
Date:   Thu Jun 24 11:55:31 2021 -0700

    RDMA/cma: Fix rdma_resolve_route() memory leak

    [ Upstream commit 74f160ead74bfe5f2b38afb4fcf86189f9ff40c9 ]

    Fix a memory leak when "mda_resolve_route() is called more than once on
```

```
    the same "rdma_cm_id".

    This is possible if cma_query_handler() triggers the
    RDMA_CM_EVENT_ROUTE_ERROR flow which puts the state machine back and
    allows rdma_resolve_route() to be called again.

    Link: https://lore.kernel.org/r/f6662b7b-bdb7-2706-1e12-47c61d3474b6@oracle.com
    Signed-off-by: Gerd Rausch <gerd.rausch@oracle.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e4d57b37f0340fbe744287ccfbe1c8f5d0538cd1
Author: Jakub Kicinski <kuba@kernel.org>
Date:   Wed Jun 23 14:44:38 2021 -0700

    net: ip: avoid OOM kills with large UDP sends over loopback

    [ Upstream commit 6d123b81ac615072a8525c13c6c41b695270a15d ]

    Dave observed number of machines hitting OOM on the UDP send
    path. The workload seems to be sending large UDP packets over
    loopback. Since loopback has MTU of 64k kernel will try to
    allocate an skb with up to 64k of head space. This has a good
    chance of failing under memory pressure. What's worse if
    the message length is <32k the allocation may trigger an
    OOM killer.

    This is entirely avoidable, we can use an skb with page frags.

    af_unix solves a similar problem by limiting the head
    length to SKB_MAX_ALLOC. This seems like a good and simple
    approach. It means that UDP messages > 16kB will now
    use fragments if underlying device supports SG, if extra
    allocator pressure causes regressions in real workloads
    we can switch to trying the large allocation first and
    falling back.

    v4: pre-calculate all the additions to alloclen so
        we can be sure it won't go over order-2

    Reported-by: Dave Jones <dsj@fb.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 602b0c1dbedc1440575b4dc1867f94d97f8c265a
Author: Martynas Pumputis <m@lambda.lt>
Date:   Wed Jun 23 15:56:45 2021 +0200

    net: retrieve netns cookie via getsocketopt

    [ Upstream commit e8b9eab99232c4e62ada9d7976c80fd5e8118289 ]

    It's getting more common to run nested container environments for
    testing cloud software. One of such examples is Kind [1] which runs a
    Kubernetes cluster in Docker containers on a single host. Each container
    acts as a Kubernetes node, and thus can run any Pod (aka container)
    inside the former. This approach simplifies testing a lot, as it
    eliminates complicated VM setups.

    Unfortunately, such a setup breaks some functionality when cgroupv2 BPF
    programs are used for load-balancing. The load-balancer BPF program
    needs to detect whether a request originates from the host netns or a
    container netns in order to allow some access, e.g. to a service via a
    loopback IP address. Typically, the programs detect this by comparing
```

```
    netns cookies with the one of the init ns via a call to
    bpf_get_netns_cookie(NULL). However, in nested environments the latter
    cannot be used given the Kubernetes node's netns is outside the init ns.
    To fix this, we need to pass the Kubernetes node netns cookie to the
    program in a different way: by extending getsockopt() with a
    SO_NETNS_COOKIE option, the orchestrator which runs in the Kubernetes
    node netns can retrieve the cookie and pass it to the program instead.

    Thus, this is following up on Eric's commit 3d368ab87cf6 ("net:
    initialize net->net_cookie at netns setup") to allow retrieval via
    SO_NETNS_COOKIE.  This is also in line in how we retrieve socket cookie
    via SO_COOKIE.

      [1] https://kind.sigs.k8s.io/

    Signed-off-by: Lorenz Bauer <lmb@cloudflare.com>
    Signed-off-by: Martynas Pumputis <m@lambda.lt>
    Cc: Eric Dumazet <edumazet@google.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b01ad7c3b08efc486e3824982858f3b30ba33c02
Author: Sean Young <sean@mess.org>
Date:   Wed Jun 23 22:37:54 2021 +0100

    media, bpf: Do not copy more entries than user space requested

    [ Upstream commit 647d446d66e493d23ca1047fa8492b0269674530 ]

    The syscall bpf(BPF_PROG_QUERY, &attr) should use the prog_cnt field to
    see how many entries user space provided and return ENOSPC if there are
    more programs than that. Before this patch, this is not checked and
    ENOSPC is never returned.

    Note that one lirc device is limited to 64 bpf programs, and user space
    I'm aware of -- ir-keytable -- always gives enough space for 64 entries
    already. However, we should not copy program ids than are requested.

    Signed-off-by: Sean Young <sean@mess.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20210623213754.632-1-sean@mess.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e6b42db7e571f5a9aa99e9d0650600b39aff9480
Author: Max Gurtovoy <mgurtovoy@nvidia.com>
Date:   Mon May 24 11:52:15 2021 +0300

    IB/isert: Align target max I/O size to initiator size

    [ Upstream commit 109d19a5eb3ddbdb87c43bfd4bcf644f4569da64 ]

    Since the Linux iser initiator default max I/O size set to 512KB and since
    there is no handshake procedure for this size in iser protocol, set the
    default max IO size of the target to 512KB as well.

    For changing the default values, there is a module parameter for both
    drivers.

    Link: https://lore.kernel.org/r/20210524085215.29005-1-mgurtovoy@nvidia.com
    Reviewed-by: Alaa Hleihel <alaa@nvidia.com>
    Reviewed-by: Israel Rukshin <israelr@nvidia.com>
    Signed-off-by: Max Gurtovoy <mgurtovoy@nvidia.com>
    Acked-by: Sagi Grimberg <sagi@grimberg.me>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
```

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

commit b3d56dd75032ea18f0126cbadf8921e6a23675e3
Author: Ilan Peer <ilan.peer@intel.com>
Date:   Fri Jun 18 13:41:37 2021 +0300

    mac80211: Properly WARN on HW scan before restart

    [ Upstream commit 45daaa1318410794de956fb8e9d06aed2dbb23d0 ]

    The following race was possible:

    1. The device driver requests HW restart.
    2. A scan is requested from user space and is propagated
       to the driver. During this flow HW_SCANNING flag is set.
    3. The thread that handles the HW restart is scheduled,
       and before starting the actual reconfiguration it
       checks that HW_SCANNING is not set. The flow does so
       without acquiring any lock, and thus the WARN fires.

    Fix this by checking that HW_SCANNING is on only after RTNL is
    acquired, i.e., user space scan request handling is no longer
    in transit.

    Signed-off-by: Ilan Peer <ilan.peer@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210618133832.8238ab3e19ab.I2693c581c70251472b4f9089e37e0
6fb2c18268f@changeid
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e32446de59d2cd10a5bd9d3bb0e7ce5ad107aad8
Author: Weilun Du <wdu@google.com>
Date:   Thu May 6 11:05:29 2021 -0700

    mac80211_hwsim: add concurrent channels scanning support over virtio

    [ Upstream commit 626c30f9e77354301ff9162c3bdddaf92d9b5cf3 ]

    This fixed the crash when setting channels to 2 or more when
    communicating over virtio.

    Signed-off-by: Weilun Du <wdu@google.com>
    Link: https://lore.kernel.org/r/20210506180530.3418576-1-wdu@google.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f9ae4a024c57a6b896f91320b8bec59c663d40c
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Thu Jun 10 22:08:16 2021 +0200

    mac80211: consider per-CPU statistics if present

    [ Upstream commit d656a4c6ead6c3f252b2f2532bc9735598f7e317 ]

    If we have been keeping per-CPU statistics, consider them
    regardless of USES_RSS, because we may not actually fill
    those, for example in non-fast-RX cases when the connection
    is not compatible with fast-RX. If we didn't fill them, the
    additional data will be zero and not affect anything, and
    if we did fill them then it's more correct to consider them.

    This fixes an issue in mesh mode where some statistics are
    not updated due to USES_RSS being set, but fast-RX isn't
```

```
    used.

    Reported-by: Thiraviyam Mariyappan <tmariyap@codeaurora.org>
    Link:
https://lore.kernel.org/r/20210610220814.13b35f5797c5.I511e9b33c5694e0d6cef4b6ae755c873d7c22
124@changeid
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5924678a442d0f5ebf33ebb76b470e68414f1318
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Wed Jun 9 15:59:44 2021 +0800

    cfg80211: fix default HE tx bitrate mask in 2G band

    [ Upstream commit 9df66d5b9f45c39b3925d16e8947cc10009b186d ]

    In 2G band, a HE sta can only supports HT and HE, but not supports VHT.
    In this case, default HE tx bitrate mask isn't filled, when we use iw to
    set bitrates without any parameter.

    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Link: https://lore.kernel.org/r/20210609075944.51130-1-pkshih@realtek.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1d88925c3b9cbe61276cfc2b1850fd795d2a9a4
Author: Gustavo A. R. Silva <gustavoars@kernel.org>
Date:   Thu Apr 22 15:00:32 2021 -0500

    wireless: wext-spy: Fix out-of-bounds warning

    [ Upstream commit e93bdd78406da9ed01554c51e38b2a02c8ef8025 ]

    Fix the following out-of-bounds warning:

    net/wireless/wext-spy.c:178:2: warning: 'memcpy' offset [25, 28] from the object at
'threshold' is out of the bounds of referenced subobject 'low' with type 'struct iw_quality'
at offset 20 [-Warray-bounds]

    The problem is that the original code is trying to copy data into a
    couple of struct members adjacent to each other in a single call to
    memcpy(). This causes a legitimate compiler warning because memcpy()
    overruns the length of &threshold.low and &spydata->spy_thr_low. As
    these are just a couple of struct members, fix this by using direct
    assignments, instead of memcpy().

    This helps with the ongoing efforts to globally enable -Warray-bounds
    and get us closer to being able to tighten the FORTIFY_SOURCE routines
    on memcpy().

    Link: https://github.com/KSPP/linux/issues/109
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Reviewed-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210422200032.GA168995@embeddedor
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 649adc384fb3544f178d66e50fc7488e146a4195
Author: Íñigo Huguet <ihuguet@redhat.com>
Date:   Mon Jun 21 17:32:36 2021 +0200

    sfc: error code if SRIOV cannot be disabled
```

```
    [ Upstream commit 1ebe4feb8b442884f5a28d2437040096723dd1ea ]

    If SRIOV cannot be disabled during device removal or module unloading,
    return error code so it can be logged properly in the calling function.

    Note that this can only happen if any VF is currently attached to a
    guest using Xen, but not with vfio/KVM. Despite that in that case the
    VFs won't work properly with PF removed and/or the module unloaded, I
    have let it as is because I don't know what side effects may have
    changing it, and also it seems to be the same that other drivers are
    doing in this situation.

    In the case of being called during SRIOV reconfiguration, the behavior
    hasn't changed because the function is called with force=false.

    Signed-off-by: Íñigo Huguet <ihuguet@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65c8d58a1f7a5574cbf0708694d500afbc845536
Author: Íñigo Huguet <ihuguet@redhat.com>
Date:   Mon Jun 21 17:32:35 2021 +0200

    sfc: avoid double pci_remove of VFs

    [ Upstream commit 45423cff1db66cf0993e8a9bd0ac93e740149e49 ]

    If pci_remove was called for a PF with VFs, the removal of the VFs was
    called twice from efx_ef10_sriov_fini: one directly with pci_driver->remove
    and another implicit by calling pci_disable_sriov, which also perform
    the VFs remove. This was leading to crashing the kernel on the second
    attempt.

    Given that pci_disable_sriov already calls to pci remove function, get
    rid of the direct call to pci_driver->remove from the driver.

    2 different ways to trigger the bug:
    - Create one or more VFs, then attach the PF to a virtual machine (at
      least with qemu/KVM)
    - Create one or more VFs, then remove the PF with:
      echo 1 > /sys/bus/pci/devices/PF_PCI_ID/remove

    Removing sfc module does not trigger the error, at least for me, because
    it removes the VF first, and then the PF.

    Example of a log with the error:
        list_del corruption, ffff967fd20a8ad0->next is LIST_POISON1 (dead000000000100)
        ------------[ cut here ]------------
        kernel BUG at lib/list_debug.c:47!
        [...trimmed...]
        RIP: 0010:__list_del_entry_valid.cold.1+0x12/0x4c
        [...trimmed...]
        Call Trace:
        efx_dissociate+0x1f/0x140 [sfc]
        efx_pci_remove+0x27/0x150 [sfc]
        pci_device_remove+0x3b/0xc0
        device_release_driver_internal+0x103/0x1f0
        pci_stop_bus_device+0x69/0x90
        pci_stop_and_remove_bus_device+0xe/0x20
        pci_iov_remove_virtfn+0xba/0x120
        sriov_disable+0x2f/0xe0
        efx_ef10_pci_sriov_disable+0x52/0x80 [sfc]
        ? pcie_aer_is_native+0x12/0x40
        efx_ef10_sriov_fini+0x72/0x110 [sfc]
        efx_pci_remove+0x62/0x150 [sfc]
```

```
         pci_device_remove+0x3b/0xc0
         device_release_driver_internal+0x103/0x1f0
         unbind_store+0xf6/0x130
         kernfs_fop_write+0x116/0x190
         vfs_write+0xa5/0x1a0
         ksys_write+0x4f/0xb0
         do_syscall_64+0x5b/0x1a0
         entry_SYSCALL_64_after_hwframe+0x65/0xca

    Signed-off-by: Íñigo Huguet <ihuguet@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 81772d2c7d043c159e024137768817ae53c8aed3
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Mon Jun 7 09:22:54 2021 +0800

    rtw88: add quirks to disable pci capabilities

    [ Upstream commit 956c6d4f20c5446727e0c912dd8f527f2dc7b779 ]

    8821CE with ASPM cannot work properly on Protempo Ltd L116HTN6SPW. Add a
    quirk to disable the cap.

    The reporter describes the symptom is that this module (driver) causes
    frequent freezes, randomly but usually within a few minutes of running
    (thus very soon after boot): screen display remains frozen, no response
    to either keyboard or mouse input. All I can do is to hold the power
    button to power off, then reboot.

    Reported-by: Paul Szabo <psz2036@gmail.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210607012254.6306-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2002b7ae2f00c0918e84a07cc9590a4d466bccf5
Author: Po-Hao Huang <phhuang@realtek.com>
Date:   Thu May 6 16:36:43 2021 +0800

    rtw88: 8822c: update RF parameter tables to v62

    [ Upstream commit 7a1baaaee6c866455c9c77bf9b0405941a3678c7 ]

    Update RTL8822C devices' RF tables to v62.
    This fixes higher than expected spur in 2400 MHz under CCK mask.

    Signed-off-by: Po-Hao Huang <phhuang@realtek.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210506083643.18317-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a7db3036585d0a03bb10c40be066f3468a04c88
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Fri Jun 18 11:01:17 2021 +0300

    iwlwifi: pcie: fix context info freeing

    [ Upstream commit 26d18c75a7496c4c52b0b6789e713dc76ebfbc87 ]

    After firmware alive, iwl_trans_pcie_gen2_fw_alive() is called
    to free the context info. However, on gen3 that will then free
    the context info with the wrong size.
```

```
    Since we free this allocation later, let it stick around until
    the device is stopped for now, freeing some of it earlier is a
    separate change.

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210618105614.afb63fb8cbc1.If4968db8e09f4ce2a1d27a6d750bc
a3d132d7d70@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b2b879f69b4da25b0097c03fd083a871dd08475
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Fri Jun 18 11:01:16 2021 +0300

    iwlwifi: pcie: free IML DMA memory allocation

    [ Upstream commit 310f60f53a86eba680d9bc20a371e13b06a5f903 ]

    In the case of gen3 devices with image loader (IML) support,
    we were leaking the IML DMA allocation and never freeing it.
    Fix that.

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210618105614.07e117dbedb7.I7bb9ebbe0617656986c2a598ea5e8
27b533bd3b9@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 44164d5528fb7c92cfbd711f8db6c6845bc42a3a
Author: Shaul Triebitz <shaul.triebitz@intel.com>
Date:   Thu Jun 17 10:08:45 2021 +0300

    iwlwifi: mvm: fix error print when session protection ends

    [ Upstream commit 976ac0af7ba2c5424bc305b926c0807d96fdcc83 ]

    When the session protection ends and the Driver is not
    associated or a beacon was not heard, the Driver
    prints "No beacons heard...".
    That's confusing for the case where not associated.
    Change the print when not associated to "Not associated...".

    Signed-off-by: Shaul Triebitz <shaul.triebitz@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210617100544.41a5a5a894fa.I9eabb76e7a3a7f4abbed8f2ef918f
1df8e825726@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 111fa98f60d205a7a6d5dd613c2c85a18ef2e1db
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Thu Jun 17 10:08:43 2021 +0300

    iwlwifi: mvm: apply RX diversity per PHY context

    [ Upstream commit a171399fd687a7d2fa56a10c9a2d7084a647677d ]

    SMPS requests may differ per interfaces due to e.g. Bluetooth
    only interfering on 2.4 GHz, so if that's the case we should,
    in the case of multiple PHY contexts, still allow RX diversity
```

```
    on PHY context that have no interfaces with SMPS requests.

    Fix the code to pass through the PHY context in question and
    skip interfaces with non-matching PHY context while iterating.

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210617100544.123c6b05809d.I992e3d1c6a29850d02eeec01712b5
b685b963a87@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89b944c896f938f5202355e1eca828705b73c363
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Sat Jun 12 14:32:38 2021 +0300

    iwlwifi: mvm: don't change band on bound PHY contexts

    [ Upstream commit 8835a64f74c46baebfc946cd5a2c861b866ebcee ]

    When we have a P2P Device active, we attempt to only change the
    PHY context it uses when we get a new remain-on-channel, if the
    P2P Device is the only user of the PHY context.

    This is fine if we're switching within a band, but if we're
    switching bands then the switch implies a removal and re-add
    of the PHY context, which isn't permitted by the firmware while
    it's bound to an interface.

    Fix the code to skip the unbind/release/... cycle only if the
    band doesn't change (or we have old devices that can switch the
    band on the fly as well.)

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20210612142637.e9ac313f70f3.I713b9d109957df7e7d9ed0861d537
7ce3f8fccd3@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit af1036530646c51c053bfb2be5ae68f32ff1377a
Author: Xiao Yang <yangx.jy@fujitsu.com>
Date:   Mon Jun 21 15:14:56 2021 +0800

    RDMA/rxe: Don't overwrite errno from ib_umem_get()

    [ Upstream commit 20ec0a6d6016aa28b9b3299be18baef1a0f91cd2 ]

    rxe_mr_init_user() always returns the fixed -EINVAL when ib_umem_get()
    fails so it's hard for user to know which actual error happens in
    ib_umem_get(). For example, ib_umem_get() will return -EOPNOTSUPP when
    trying to pin pages on a DAX file.

    Return actual error as mlx4/mlx5 does.

    Link: https://lore.kernel.org/r/20210621071456.4259-1-ice_yangxiao@163.com
    Signed-off-by: Xiao Yang <yangx.jy@fujitsu.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e34e09aa418117fca3eb7d14d748e65a59da4ca
Author: Logush Oliver <ollogush@amd.com>
Date:   Tue Jun 8 15:27:24 2021 -0400
```

```
    drm/amd/display: Fix edp_bootup_bl_level initialization issue

    [ Upstream commit eeb90e26ed05dd44553d557057bf35f08f853af8 ]

    [why]
    Updating the file to fix the missing line

    Signed-off-by: Logush Oliver <ollogush@amd.com>
    Reviewed-by: Charlene Liu <Charlene.Liu@amd.com>
    Acked-by: Bindu Ramamurthy <bindu.r@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 f033a30e32cefb956cf0a18afca17c8a267aa147
Author: Longpeng(Mike) <longpeng2@huawei.com>
Date:   Mon Jun 21 14:26:01 2021 +0800

    vsock: notify server to shutdown when client has pending signal

    [ Upstream commit c7ff9cff70601ea19245d997bb977344663434c7 ]

    The client's sk_state will be set to TCP_ESTABLISHED if the server
    replay the client's connect request.

    However, if the client has pending signal, its sk_state will be set
    to TCP_CLOSE without notify the server, so the server will hold the
    corrupt connection.

                client                          server

    1. sk_state=TCP_SYN_SENT         |
    2. call ->connect()              |
    3. wait reply                    |
                                     | 4. sk_state=TCP_ESTABLISHED
                                     | 5. insert to connected list
                                     | 6. reply to the client
    7. sk_state=TCP_ESTABLISHED      |
    8. insert to connected list      |
    9. *signal pending* <---------------------- the user kill client
    10. sk_state=TCP_CLOSE           |
    client is exiting...             |
    11. call ->release()             |
         virtio_transport_close
          if (!(sk->sk_state == TCP_ESTABLISHED ||
                  sk->sk_state == TCP_CLOSING))
                    return true; *return at here, the server cannot notice the connection is
corrupt*

    So the client should notify the peer in this case.

    Cc: David S. Miller <davem@davemloft.net>
    Cc: Jakub Kicinski <kuba@kernel.org>
    Cc: Jorgen Hansen <jhansen@vmware.com>
    Cc: Norbert Slusarek <nslusarek@gmx.net>
    Cc: Andra Paraschiv <andraprs@amazon.com>
    Cc: Colin Ian King <colin.king@canonical.com>
    Cc: David Brazdil <dbrazdil@google.com>
    Cc: Alexander Popov <alex.popov@linux.com>
    Suggested-by: Stefano Garzarella <sgarzare@redhat.com>
    Link: https://lkml.org/lkml/2021/5/17/418
    Signed-off-by: lixianming <lixianming5@huawei.com>
    Signed-off-by: Longpeng(Mike) <longpeng2@huawei.com>
    Reviewed-by: Stefano Garzarella <sgarzare@redhat.com>
```

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

commit ac05406330eb17b2dd70e48011d374831388c0ab
Author: Zheyu Ma <zheyuma97@gmail.com>
Date:   Sun Jun 20 15:24:15 2021 +0000

    atm: nicstar: register the interrupt handler in the right place

    [ Upstream commit 70b639dc41ad499384e41e106fce72e36805c9f2 ]

    Because the error handling is sequential, the application of resources
    should be carried out in the order of error handling, so the operation
    of registering the interrupt handler should be put in front, so as not
    to free the unregistered interrupt handler during error handling.

    This log reveals it:

    [    3.438724] Trying to free already-free IRQ 23
    [    3.439060] WARNING: CPU: 5 PID: 1 at kernel/irq/manage.c:1825 free_irq+0xfb/0x480
    [    3.440039] Modules linked in:
    [    3.440257] CPU: 5 PID: 1 Comm: swapper/0 Not tainted 5.12.4-g70e7f0549188-dirty #142
    [    3.440793] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.12.0-59-
gc9ba5276e321-prebuilt.qemu.org 04/01/2014
    [    3.441561] RIP: 0010:free_irq+0xfb/0x480
    [    3.441845] Code: 6e 08 74 6f 4d 89 f4 e8 c3 78 09 00 4d 8b 74 24 18 4d 85 f6 75 e3
e8 b4 78 09 00 8b 75 c8 48 c7 c7 a0 ac d5 85 e8 95 d7 f5 ff <0f> 0b 48 8b 75 c0 4c 89 ff e8
87 c5 90 03 48 8b 43 40 4c 8b a0 80
    [    3.443121] RSP: 0000:ffffc90000017b50 EFLAGS: 00010086
    [    3.443483] RAX: 0000000000000000 RBX: ffff888107c6f000 RCX: 0000000000000000
    [    3.443972] RDX: 0000000000000000 RSI: ffffffff8123f301 RDI: 00000000ffffffff
    [    3.444462] RBP: ffffc90000017b90 R08: 0000000000000001 R09: 0000000000000003
    [    3.444950] R10: 0000000000000000 R11: 0000000000000001 R12: 0000000000000000
    [    3.444994] R13: ffff888107dc0000 R14: ffff888104f6bf00 R15: ffff888107c6f0a8
    [    3.444994] FS:  0000000000000000(0000) GS:ffff88817bd40000(0000)
knlGS:0000000000000000
    [    3.444994] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [    3.444994] CR2: 0000000000000000 CR3: 000000000642e000 CR4: 00000000000006e0
    [    3.444994] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [    3.444994] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [    3.444994] Call Trace:
    [    3.444994]  ns_init_card_error+0x18e/0x250
    [    3.444994]  nicstar_init_one+0x10d2/0x1130
    [    3.444994]  local_pci_probe+0x4a/0xb0
    [    3.444994]  pci_device_probe+0x126/0x1d0
    [    3.444994]  ? pci_device_remove+0x100/0x100
    [    3.444994]  really_probe+0x27e/0x650
    [    3.444994]  driver_probe_device+0x84/0x1d0
    [    3.444994]  ? mutex_lock_nested+0x16/0x20
    [    3.444994]  device_driver_attach+0x63/0x70
    [    3.444994]  __driver_attach+0x117/0x1a0
    [    3.444994]  ? device_driver_attach+0x70/0x70
    [    3.444994]  bus_for_each_dev+0xb6/0x110
    [    3.444994]  ? rdinit_setup+0x40/0x40
    [    3.444994]  driver_attach+0x22/0x30
    [    3.444994]  bus_add_driver+0x1e6/0x2a0
    [    3.444994]  driver_register+0xa4/0x180
    [    3.444994]  __pci_register_driver+0x77/0x80
    [    3.444994]  ? uPD98402_module_init+0xd/0xd
    [    3.444994]  nicstar_init+0x1f/0x75
    [    3.444994]  do_one_initcall+0x7a/0x3d0
    [    3.444994]  ? rdinit_setup+0x40/0x40
    [    3.444994]  ? rcu_read_lock_sched_held+0x4a/0x70
    [    3.444994]  kernel_init_freeable+0x2a7/0x2f9
    [    3.444994]  ? rest_init+0x2c0/0x2c0
```

```
     [    3.444994]  kernel_init+0x13/0x180
     [    3.444994]  ? rest_init+0x2c0/0x2c0
     [    3.444994]  ? rest_init+0x2c0/0x2c0
     [    3.444994]  ret_from_fork+0x1f/0x30
     [    3.444994] Kernel panic - not syncing: panic_on_warn set ...
     [    3.444994] CPU: 5 PID: 1 Comm: swapper/0 Not tainted 5.12.4-g70e7f0549188-dirty #142
     [    3.444994] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.12.0-59-
  gc9ba5276e321-prebuilt.qemu.org 04/01/2014
     [    3.444994] Call Trace:
     [    3.444994]  dump_stack+0xba/0xf5
     [    3.444994]  ? free_irq+0xfb/0x480
     [    3.444994]  panic+0x155/0x3ed
     [    3.444994]  ? __warn+0xed/0x150
     [    3.444994]  ? free_irq+0xfb/0x480
     [    3.444994]  __warn+0x103/0x150
     [    3.444994]  ? free_irq+0xfb/0x480
     [    3.444994]  report_bug+0x119/0x1c0
     [    3.444994]  handle_bug+0x3b/0x80
     [    3.444994]  exc_invalid_op+0x18/0x70
     [    3.444994]  asm_exc_invalid_op+0x12/0x20
     [    3.444994] RIP: 0010:free_irq+0xfb/0x480
     [    3.444994] Code: 6e 08 74 6f 4d 89 f4 e8 c3 78 09 00 4d 8b 74 24 18 4d 85 f6 75 e3
  e8 b4 78 09 00 8b 75 c8 48 c7 c7 a0 ac d5 85 e8 95 d7 f5 ff <0f> 0b 48 8b 75 c0 4c 89 ff e8
  87 c5 90 03 48 8b 43 40 4c 8b a0 80
     [    3.444994] RSP: 0000:ffffc90000017b50 EFLAGS: 00010086
     [    3.444994] RAX: 0000000000000000 RBX: ffff888107c6f000 RCX: 0000000000000000
     [    3.444994] RDX: 0000000000000000 RSI: ffffffff8123f301 RDI: 00000000ffffffff
     [    3.444994] RBP: ffffc90000017b90 R08: 0000000000000001 R09: 0000000000000003
     [    3.444994] R10: 0000000000000000 R11: 0000000000000001 R12: 0000000000000000
     [    3.444994] R13: ffff888107dc0000 R14: ffff888104f6bf00 R15: ffff888107c6f0a8
     [    3.444994]  ? vprintk_func+0x71/0x110
     [    3.444994]  ns_init_card_error+0x18e/0x250
     [    3.444994]  nicstar_init_one+0x10d2/0x1130
     [    3.444994]  local_pci_probe+0x4a/0xb0
     [    3.444994]  pci_device_probe+0x126/0x1d0
     [    3.444994]  ? pci_device_remove+0x100/0x100
     [    3.444994]  really_probe+0x27e/0x650
     [    3.444994]  driver_probe_device+0x84/0x1d0
     [    3.444994]  ? mutex_lock_nested+0x16/0x20
     [    3.444994]  device_driver_attach+0x63/0x70
     [    3.444994]  __driver_attach+0x117/0x1a0
     [    3.444994]  ? device_driver_attach+0x70/0x70
     [    3.444994]  bus_for_each_dev+0xb6/0x110
     [    3.444994]  ? rdinit_setup+0x40/0x40
     [    3.444994]  driver_attach+0x22/0x30
     [    3.444994]  bus_add_driver+0x1e6/0x2a0
     [    3.444994]  driver_register+0xa4/0x180
     [    3.444994]  __pci_register_driver+0x77/0x80
     [    3.444994]  ? uPD98402_module_init+0xd/0xd
     [    3.444994]  nicstar_init+0x1f/0x75
     [    3.444994]  do_one_initcall+0x7a/0x3d0
     [    3.444994]  ? rdinit_setup+0x40/0x40
     [    3.444994]  ? rcu_read_lock_sched_held+0x4a/0x70
     [    3.444994]  kernel_init_freeable+0x2a7/0x2f9
     [    3.444994]  ? rest_init+0x2c0/0x2c0
     [    3.444994]  kernel_init+0x13/0x180
     [    3.444994]  ? rest_init+0x2c0/0x2c0
     [    3.444994]  ? rest_init+0x2c0/0x2c0
     [    3.444994]  ret_from_fork+0x1f/0x30
     [    3.444994] Dumping ftrace buffer:
     [    3.444994]    (ftrace buffer empty)
     [    3.444994] Kernel Offset: disabled
     [    3.444994] Rebooting in 1 seconds..
```

Signed-off-by: Zheyu Ma <zheyuma97@gmail.com>

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

commit 3c0b0b2533f9fd91aa10dfdd48bc27f3526e585c
Author: Zheyu Ma <zheyuma97@gmail.com>
Date:   Sun Jun 20 15:24:14 2021 +0000

    atm: nicstar: use 'dma_free_coherent' instead of 'kfree'

    [ Upstream commit 6a1e5a4af17e440dd82a58a2c5f40ff17a82b722 ]

    When 'nicstar_init_one' fails, 'ns_init_card_error' will be executed for
    error handling, but the correct memory free function should be used,
    otherwise it will cause an error. Since 'card->rsq.org' and
    'card->tsq.org' are allocated using 'dma_alloc_coherent' function, they
    should be freed using 'dma_free_coherent'.

    Fix this by using 'dma_free_coherent' instead of 'kfree'

    This log reveals it:

    [    3.440294] kernel BUG at mm/slub.c:4206!
    [    3.441059] invalid opcode: 0000 [#1] PREEMPT SMP PTI
    [    3.441430] CPU: 2 PID: 1 Comm: swapper/0 Not tainted 5.12.4-g70e7f0549188-dirty #141
    [    3.441986] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.12.0-59-
gc9ba5276e321-prebuilt.qemu.org 04/01/2014
    [    3.442780] RIP: 0010:kfree+0x26a/0x300
    [    3.443065] Code: e8 3a c3 b9 ff e9 d6 fd ff ff 49 8b 45 00 31 db a9 00 00 01 00 75
4d 49 8b 45 00 a9 00 00 01 00 75 0a 49 8b 45 08 a8 01 75 02 <0f> 0b 89 d9 b8 00 10 00 00 be
06 00 00 00 48 d3 e0 f7 d8 48 63 d0
    [    3.443396] RSP: 0000:ffffc90000017b70 EFLAGS: 00010246
    [    3.443396] RAX: dead000000000100 RBX: 0000000000000000 RCX: 0000000000000000
    [    3.443396] RDX: 0000000000000000 RSI: ffffffff85d3df94 RDI: ffffffff85df38e6
    [    3.443396] RBP: ffffc90000017b90 R08: 0000000000000001 R09: 0000000000000001
    [    3.443396] R10: 0000000000000000 R11: 0000000000000001 R12: ffff888107dc0000
    [    3.443396] R13: ffffea00001f0100 R14: ffff888101a8bf00 R15: ffff888107dc0160
    [    3.443396] FS:  0000000000000000(0000) GS:ffff88817bc80000(0000)
knlGS:0000000000000000
    [    3.443396] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [    3.443396] CR2: 0000000000000000 CR3: 000000000642e000 CR4: 00000000000006e0
    [    3.443396] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [    3.443396] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [    3.443396] Call Trace:
    [    3.443396]  ns_init_card_error+0x12c/0x220
    [    3.443396]  nicstar_init_one+0x10d2/0x1130
    [    3.443396]  local_pci_probe+0x4a/0xb0
    [    3.443396]  pci_device_probe+0x126/0x1d0
    [    3.443396]  ? pci_device_remove+0x100/0x100
    [    3.443396]  really_probe+0x27e/0x650
    [    3.443396]  driver_probe_device+0x84/0x1d0
    [    3.443396]  ? mutex_lock_nested+0x16/0x20
    [    3.443396]  device_driver_attach+0x63/0x70
    [    3.443396]  __driver_attach+0x117/0x1a0
    [    3.443396]  ? device_driver_attach+0x70/0x70
    [    3.443396]  bus_for_each_dev+0xb6/0x110
    [    3.443396]  ? rdinit_setup+0x40/0x40
    [    3.443396]  driver_attach+0x22/0x30
    [    3.443396]  bus_add_driver+0x1e6/0x2a0
    [    3.443396]  driver_register+0xa4/0x180
    [    3.443396]  __pci_register_driver+0x77/0x80
    [    3.443396]  ? uPD98402_module_init+0xd/0xd
    [    3.443396]  nicstar_init+0x1f/0x75
    [    3.443396]  do_one_initcall+0x7a/0x3d0
    [    3.443396]  ? rdinit_setup+0x40/0x40
    [    3.443396]  ? rcu_read_lock_sched_held+0x4a/0x70
```

```
     [    3.443396]  kernel_init_freeable+0x2a7/0x2f9
     [    3.443396]  ? rest_init+0x2c0/0x2c0
     [    3.443396]  kernel_init+0x13/0x180
     [    3.443396]  ? rest_init+0x2c0/0x2c0
     [    3.443396]  ? rest_init+0x2c0/0x2c0
     [    3.443396]  ret_from_fork+0x1f/0x30
     [    3.443396] Modules linked in:
     [    3.443396] Dumping ftrace buffer:
     [    3.443396]    (ftrace buffer empty)
     [    3.458593] ---[ end trace 3c6f8f0d8ef59bcd ]---
     [    3.458922] RIP: 0010:kfree+0x26a/0x300
     [    3.459198] Code: e8 3a c3 b9 ff e9 d6 fd ff ff 49 8b 45 00 31 db a9 00 00 01 00 75
 4d 49 8b 45 00 a9 00 00 01 00 75 0a 49 8b 45 08 a8 01 75 02 <0f> 0b 89 d9 b8 00 10 00 00 be
 06 00 00 00 48 d3 e0 f7 d8 48 63 d0
     [    3.460499] RSP: 0000:ffffc90000017b70 EFLAGS: 00010246
     [    3.460870] RAX: dead000000000100 RBX: 0000000000000000 RCX: 0000000000000000
     [    3.461371] RDX: 0000000000000000 RSI: ffffffff85d3df94 RDI: ffffffff85df38e6
     [    3.461873] RBP: ffffc90000017b90 R08: 0000000000000001 R09: 0000000000000001
     [    3.462372] R10: 0000000000000000 R11: 0000000000000001 R12: ffff888107dc0000
     [    3.462871] R13: ffffea00001f0100 R14: ffff888101a8bf00 R15: ffff888107dc0160
     [    3.463368] FS:  0000000000000000(0000) GS:ffff88817bc80000(0000)
 knlGS:0000000000000000
     [    3.463949] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
     [    3.464356] CR2: 0000000000000000 CR3: 000000000642e000 CR4: 00000000000006e0
     [    3.464856] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
     [    3.465356] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
     [    3.465860] Kernel panic - not syncing: Fatal exception
     [    3.466370] Dumping ftrace buffer:
     [    3.466616]    (ftrace buffer empty)
     [    3.466871] Kernel Offset: disabled
     [    3.467122] Rebooting in 1 seconds..

     Signed-off-by: Zheyu Ma <zheyuma97@gmail.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 5ad7f389f93b5227f2f02e29deb88f2ee53b8561
 Author: Fugang Duan <fugang.duan@nxp.com>
 Date:   Mon Jun 21 14:27:37 2021 +0800

     net: fec: add ndo_select_queue to fix TX bandwidth fluctuations

     [ Upstream commit 52c4a1a85f4b346c39c896c0168f4a843b3385ff ]

     As we know that AVB is enabled by default, and the ENET IP design is
     queue 0 for best effort, queue 1&2 for AVB Class A&B. Bandwidth of each
     queue 1&2 set in driver is 50%, TX bandwidth fluctuated when selecting
     tx queues randomly with FEC_QUIRK_HAS_AVB quirk available.

     This patch adds ndo_select_queue callback to select queues for
     transmitting to fix this issue. It will always return queue 0 if this is
     not a vlan packet, and return queue 1 or 2 based on priority of vlan
     packet.

     You may complain that in fact we only use single queue for trasmitting
     if we are not targeted to VLAN. Yes, but seems we have no choice, since
     AVB is enabled when the driver probed, we can't switch this feature
     dynamicly. After compare multiple queues to single queue, TX throughput
     almost no improvement.

     One way we can implemet is to configure the driver to multiple queues
     with Round-robin scheme by default. Then add ndo_setup_tc callback to
     enable/disable AVB feature for users. Unfortunately, ENET AVB IP seems
     not follow the standard 802.1Qav spec. We only can program
     DMAnCFG[IDLE_SLOPE] field to calculate bandwidth fraction. And idle
```

```
    slope is restricted to certain valus (a total of 19). It's far away from
    CBS QDisc implemented in Linux TC framework. If you strongly suggest to do
    this, I think we only can support limited numbers of bandwidth and reject
    others, but it's really urgly and wried.

    With this patch, VLAN tagged packets route to queue 0/1/2 based on vlan
    priority; VLAN untagged packets route to queue 0.

    Tested-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Reported-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Signed-off-by: Fugang Duan <fugang.duan@nxp.com>
    Signed-off-by: Joakim Zhang <qiangqing.zhang@nxp.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df628a2027551dc339b3b001fa47a2cf9c292cf1
Author: Joakim Zhang <qiangqing.zhang@nxp.com>
Date:   Mon Jun 21 14:27:36 2021 +0800

    net: fec: add FEC_QUIRK_HAS_MULTI_QUEUES represents i.MX6SX ENET IP

    [ Upstream commit 471ff4455d61c9929ae912328859921708e1eafc ]

    Frieder Schrempf reported a TX throuthput issue [1], it happens quite often
    that the measured bandwidth in TX direction drops from its expected/nominal
    value to something like ~50% (for 100M) or ~67% (for 1G) connections.

    [1] https://lore.kernel.org/linux-arm-kernel/421cc86c-b66f-b372-32f7-
21e59f9a98bc@kontron.de/

    The issue becomes clear after digging into it, Net core would select
    queues when transmitting packets. Since FEC have not impletemented
    ndo_select_queue callback yet, so it will call netdev_pick_tx to select
    queues randomly.

    For i.MX6SX ENET IP with AVB support, driver default enables this
    feature. According to the setting of QOS/RCMRn/DMAnCFG registers, AVB
    configured to Credit-based scheme, 50% bandwidth of each queue 1&2.

    With below tests let me think more:
    1) With FEC_QUIRK_HAS_AVB quirk, can reproduce TX bandwidth fluctuations issue.
    2) Without FEC_QUIRK_HAS_AVB quirk, can't reproduce TX bandwidth fluctuations issue.

    The related difference with or w/o FEC_QUIRK_HAS_AVB quirk is that, whether we
    program FTYPE field of TxBD or not. As I describe above, AVB feature is
    enabled by default. With FEC_QUIRK_HAS_AVB quirk, frames in queue 0
    marked as non-AVB, and frames in queue 1&2 marked as AVB Class A&B. It's
    unreasonable if frames in queue 1&2 are not required to be time-sensitive.
    So when Net core select tx queues ramdomly, Credit-based scheme would work
    and lead to TX bandwidth fluctuated. On the other hand, w/o
    FEC_QUIRK_HAS_AVB quirk, frames in queue 1&2 are all marked as non-AVB, so
    Credit-based scheme would not work.

    Till now, how can we fix this TX throughput issue? Yes, please remove
    FEC_QUIRK_HAS_AVB quirk if you suffer it from time-nonsensitive networking.
    However, this quirk is used to indicate i.MX6SX, other setting depends
    on it. So this patch adds a new quirk FEC_QUIRK_HAS_MULTI_QUEUES to
    represent i.MX6SX, it is safe for us remove FEC_QUIRK_HAS_AVB quirk
    now.

    FEC_QUIRK_HAS_AVB quirk is set by default in the driver, and users may
    not know much about driver details, they would waste effort to find the
    root cause, that is not we want. The following patch is a implementation
    to fix it and users don't need to modify the driver.
```

```
    Tested-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Reported-by: Frieder Schrempf <frieder.schrempf@kontron.de>
    Signed-off-by: Joakim Zhang <qiangqing.zhang@nxp.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b223fca477adfa86db681c9af1f23bfee9f98494
Author: Huang Pei <huangpei@loongson.cn>
Date:   Fri Jun 11 15:09:46 2021 +0800

    MIPS: add PMD table accounting into MIPS'pmd_alloc_one

    [ Upstream commit ed914d48b6a1040d1039d371b56273d422c0081e ]

    This fixes Page Table accounting bug.

    MIPS is the ONLY arch just defining __HAVE_ARCH_PMD_ALLOC_ONE alone.
    Since commit b2b29d6d011944 (mm: account PMD tables like PTE tables),
    "pmd_free" in asm-generic with PMD table accounting and "pmd_alloc_one"
    in MIPS without PMD table accounting causes PageTable accounting number
    negative, which read by global_zone_page_state(), always returns 0.

    Signed-off-by: Huang Pei <huangpei@loongson.cn>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59dc71f577eac2e4aa56c809fc3d827539bcc8a8
Author: Pascal Terjan <pterjan@google.com>
Date:   Sat Apr 24 18:29:59 2021 +0100

    rtl8xxxu: Fix device info for RTL8192EU devices

    [ Upstream commit c240b044edefa3c3af4014a4030e017dd95b59a1 ]

    Based on 2001:3319 and 2357:0109 which I used to test the fix and
    0bda:818b and 2357:0108 for which I found efuse dumps online.

    == 2357:0109 ==
    === Before ===
    Vendor: Realtek
    Product: \x03802.11n NI
    Serial:
    === After ===
    Vendor: Realtek
    Product: 802.11n NIC
    Serial not available.

    == 2001:3319 ==
    === Before ===
    Vendor: Realtek
    Product: Wireless N
    Serial: no USB Adap
    === After ===
    Vendor: Realtek
    Product: Wireless N Nano USB Adapter
    Serial not available.

    Signed-off-by: Pascal Terjan <pterjan@google.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210424172959.1559890-1-pterjan@google.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97e50ed9eae588cb274fa499375a1d49ca24cfa9
Author: Ryder Lee <ryder.lee@mediatek.com>
```

```
Date:   Thu Jun 17 15:17:49 2021 +0800

    mt76: fix iv and CCMP header insertion

    [ Upstream commit c368362c36d3d4cedbc9a1c9caa95960912cc429 ]

    The iv from RXD is only for TKIP_RSC/CCMP_PN/GCMP_PN, and it needs a
    check for CCMP header insertion. Move mt76_cipher_type to mt76.h to
    reduce duplicated code.

    Signed-off-by: Xing Song <xing.song@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65c015b505159e169f81482466b747b306489ade
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Sat Jun 12 02:04:20 2021 +0800

    mt76: mt7915: fix IEEE80211_HE_PHY_CAP7_MAX_NC for station mode

    [ Upstream commit 2707ff4dd7b1479dbd44ebb3c74788084cc95245 ]

    The value of station mode is always 0.

    Fixed: 00b2e16e0063 ("mt76: mt7915: add TxBF capabilities")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b6248b3e6498ce0985b9f9372123de35c362be1
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Wed Jun 9 14:15:32 2021 +0800

    mt76: connac: fix the maximum interval schedule scan can support

    [ Upstream commit abded041a07467c2f3dfe10afd9ea10572c63cc9 ]

    Maximum interval (in seconds) for schedule scan plan supported by
    the offload firmware can be U16_MAX.

    Signed-off-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbe1d0ada80eeb8ff745c2780b7bbcc0fcb3877f
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Wed Jun 2 18:00:14 2021 +0200

    mt76: connac: fix UC entry is being overwritten

    [ Upstream commit 82453b1cbf9ef166364c12b5464251f16bac5f51 ]

    Fix UC entry is being overwritten by BC entry

    Tested-by: Deren Wu <deren.wu@mediatek.com>
    Co-developed-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57602135190b0d9324aa82e3cc01cdbfaf430ffe
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Mon May 31 08:33:18 2021 +0200
```

```
    mt76: mt7921: enable hw offloading for wep keys

    [ Upstream commit a60951d4faa0ef2e475797dd217c2eaee32ed1c2 ]

    Enable wep key hw offloading for sta mode. This patch fixes
    WoW support for wep connections.

    Tested-by: Deren.Wu <deren.wu@mediatek.com>
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 869417fa825c82c21e35742d72d0bbc27525263f
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Sun May 23 23:10:12 2021 +0200

    mt76: mt7921: reset wfsys during hw probe

    [ Upstream commit 01f7da40917923bf9d8fd8d5c9a6ed646004e47c ]

    This patch fixes a mcu hang during device probe on
    Marvell ESPRESSObin after a hot reboot.

    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6620a41fe435b42e5c869cced265bef6f0ad94d2
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Thu May 20 11:46:37 2021 +0800

    mt76: mt7921: fix reset under the deep sleep is enabled

    [ Upstream commit 7bf0a71e839822bb6ba04a6e163ad334314e2659 ]

    To fix possibly the race to access register between the WiFi reset
    and the other context that is caused by explicitly cancelling ps_work
    and wake_work to break PM_STATE consistency.

    Deep sleep would cause the hardware into the inactive state,
    so we forcely put device drv_own state before we start to reset.

    The patch also ignore the reset request when the procedure is in
    progress to avoid the consecutive WiFi resets.

    localhost ~ # [ 2932.073966] SError Interrupt on CPU7, code 0xbe000011
    [ 2932.073967] CPU: 7 PID: 8761 Comm: kworker/u16:2 Not tainted 5.4.112 #30
    [ 2932.073968] Hardware name: MediaTek Asurada rev1 board (DT)
    [ 2932.073968] Workqueue: phy0 ieee80211_reconfig_filter [mac80211]
    [ 2932.073969] pstate: 80400089 (Nzcv daIf +PAN -UAO)
    [ 2932.073969] pc : el1_irq+0x78/0x180
    [ 2932.073970] lr : mt76_mmio_rmw+0x30/0x5c [mt76]
    [ 2932.073970] sp : ffffffc01142bad0
    [ 2932.073970] x29: ffffffc01142bc00 x28: ffffff8f96fb1e00
    [ 2932.073971] x27: ffffffd2cdc12138 x26: ffffffd2cdaeb018
    [ 2932.073972] x25: 0000000000000000 x24: ffffff8fa8e14c08
    [ 2932.073973] x23: 0000000080c00009 x22: ffffffd2a5603918
    [ 2932.073974] x21: ffffffc01142bc10 x20: 0000007fffffffff
    [ 2932.073975] x19: 0000000000000000 x18: 0000000000000400
    [ 2932.073975] x17: 0000000000000400 x16: ffffffd2cd2b87dc
    [ 2932.073976] x15: 0000000000000000 x14: 0000000000000000
    [ 2932.073977] x13: 0000000000000001 x12: 0000000000000001
    [ 2932.073978] x11: 0000000000000001 x10: 000000000010e000
    [ 2932.073978] x9 : 0000000000000000 x8 : ffffffc013921404
```

```
     [ 2932.073979] x7 : 000000b2b5593519 x6 : 0000000000300000
     [ 2932.073980] x5 : 0000000000000000 x4 : ffffffc01142bbc8
     [ 2932.073980] x3 : 00000000000001f0 x2 : 0000000000000000
     [ 2932.073981] x1 : 0000000000021404 x0 : ffffff8fa8e12300
     [ 2932.073982] Kernel panic - not syncing: Asynchronous SError Interrupt
     [ 2932.073983] CPU: 7 PID: 8761 Comm: kworker/u16:2 Not tainted 5.4.112 #30
     [ 2932.073983] Hardware name: MediaTek Asurada rev1 board (DT)
     [ 2932.073984] Workqueue: phy0 ieee80211_reconfig_filter [mac80211]
     [ 2932.073984] Call trace:
     [ 2932.073985]  dump_backtrace+0x0/0x14c
     [ 2932.073985]  show_stack+0x20/0x2c
     [ 2932.073985]  dump_stack+0xa0/0xf8
     [ 2932.073986]  panic+0x154/0x360
     [ 2932.073986]  test_taint+0x0/0x44
     [ 2932.073986]  arm64_serror_panic+0x78/0x84
     [ 2932.073987]  do_serror+0x0/0x118
     [ 2932.073987]  do_serror+0xa4/0x118
     [ 2932.073987]  el1_error+0x84/0xf8
     [ 2932.073988]  el1_irq+0x78/0x180
     [ 2932.073988]  mt76_mmio_rr+0x30/0xf0 [mt76]
     [ 2932.073988]  mt76_mmio_rmw+0x30/0x5c [mt76]
     [ 2932.073989]  mt7921_rmw+0x4c/0x5c [mt7921e]
     [ 2932.073989]  mt7921_configure_filter+0x138/0x160 [mt7921e]
     [ 2932.073990]  ieee80211_configure_filter+0x2f0/0x3e0 [mac80211]
     [ 2932.073990]  ieee80211_reconfig_filter+0x1c/0x28 [mac80211]
     [ 2932.073990]  process_one_work+0x208/0x3c8
     [ 2932.073991]  worker_thread+0x23c/0x3e8
     [ 2932.073991]  kthread+0x140/0x17c
     [ 2932.073992]  ret_from_fork+0x10/0x18
     [ 2932.074071] SMP: stopping secondary CPUs
     [ 2932.074071] Kernel Offset: 0x12bc800000 from 0xffffffc010000000
     [ 2932.074072] PHYS_OFFSET: 0xffffffff80000000
     [ 2932.074072] CPU features: 0x080026,2a80aa18
     [ 2932.074072] Memory Limit: none

     Co-developed-by: Lorenzo Bianconi <lorenzo@kernel.org>
     Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
     Signed-off-by: Sean Wang <sean.wang@mediatek.com>
     Signed-off-by: Felix Fietkau <nbd@nbd.name>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 7922a844969d6238e8dadab8608b61b179087483
  Author: Evelyn Tsai <evelyn.tsai@mediatek.com>
  Date:   Mon May 3 16:04:37 2021 +0800

     mt76: mt7915: fix tssi indication field of DBDC NICs

     [ Upstream commit 64cf5ad3c2fa841e4b416343a7ea69c63d60fa4e ]

     Correct the bitfield which indicates TSSI on/off for MT7915D NIC.

     Signed-off-by: Evelyn Tsai <evelyn.tsai@mediatek.com>
     Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
     Signed-off-by: Felix Fietkau <nbd@nbd.name>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 38b73fe18c6b0628bd4156f13775c484f5c650db
  Author: xinhui pan <xinhui.pan@amd.com>
  Date:   Tue Jun 15 15:11:07 2021 +0800

     drm/amdkfd: Walk through list with dqm lock hold

     [ Upstream commit 56f221b6389e7ab99c30bbf01c71998ae92fc584 ]

     To avoid any list corruption.
```

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

commit af15052971e17de15f3817b9918e9f46d29c67d5
Author: Stanley.Yang <Stanley.Yang@amd.com>
Date:   Wed Jun 16 21:14:01 2021 +0800

    drm/amdgpu: fix bad address translation for sienna_cichlid

    [ Upstream commit 6ec598cc9dfbf40433e94a2ed1a622e3ef80268b ]

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

commit a8eca6968b1e66785ef142db65afa62c26ec1029
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Thu Jun 17 18:13:59 2021 +0100

    io_uring: fix false WARN_ONCE

    [ Upstream commit e6ab8991c5d0b0deae0961dc22c0edd1dee328f5 ]

    WARNING: CPU: 1 PID: 11749 at fs/io-wq.c:244 io_wqe_wake_worker fs/io-wq.c:244 [inline]
    WARNING: CPU: 1 PID: 11749 at fs/io-wq.c:244 io_wqe_enqueue+0x7f6/0x910 fs/io-wq.c:751

    A WARN_ON_ONCE() in io_wqe_wake_worker() can be triggered by a valid
    userspace setup. Replace it with pr_warn.

    Reported-by: syzbot+ea2f1484cffe5109dc10@syzkaller.appspotmail.com
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/f7ede342c3342c4c26668f5168e2993e38bbd99c.1623949695.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9e134247013e29e0a0c13365f47f1394824a522f
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Jun 17 16:02:07 2021 +0800

    net: sched: fix error return code in tcf_del_walker()

    [ Upstream commit 55d96f72e8ddc0a294e0b9c94016edbb699537e1 ]

    When nla_put_u32() fails, 'ret' could be 0, it should
    return error code in tcf_del_walker().

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06648197e4ae90ef843ad1656ae22108be26a944
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Jun 17 13:11:19 2021 +0800

    net: ipa: Add missing of_node_put() in ipa_firmware_load()

    [ Upstream commit b244163f2c45c12053cb0291c955f892e79ed8a9 ]
```

```
    This node pointer is returned by of_parse_phandle() with refcount
    incremented in this function. of_node_put() on it before exiting
    this function.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Alex Elder <elder@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0147b9c4ebf6259638a3a590a5c1f1919af3ee00
Author: Jian Shen <shenjian15@huawei.com>
Date:   Thu Jun 17 11:37:11 2021 +0800

    net: fix mistake path for netdev_features_strings

    [ Upstream commit 2d8ea148e553e1dd4e80a87741abdfb229e2b323 ]

    Th_strings arrays netdev_features_strings, tunable_strings, and
    phy_tunable_strings has been moved to file net/ethtool/common.c.
    So fixes the comment.

    Signed-off-by: Jian Shen <shenjian15@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d23e6082729e472b7182ca9ed0ccc6e94863bb9
Author: Felix Fietkau <nbd@nbd.name>
Date:   Sun May 9 20:23:01 2021 +0200

    mt76: dma: use ieee80211_tx_status_ext to free packets when tx fails

    [ Upstream commit 94e4f5794627a80ce036c35b32a9900daeb31be3 ]

    Fixes AQL issues on full queues, especially with 802.3 encap offload

    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2a64f3732310a16e746650e49f55ba5dc71ca59a
Author: Felix Fietkau <nbd@nbd.name>
Date:   Fri May 7 14:07:53 2021 +0200

    mt76: mt7615: fix fixed-rate tx status reporting

    [ Upstream commit ec8f1a90d006f7cedcf86ef19fd034a406a213d6 ]

    Rely on the txs fixed-rate bit instead of info->control.rates

    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2b69554431f44dcfa51a57f7dbde3461a435747
Author: Jacob Keller <jacob.e.keller@intel.com>
Date:   Thu May 6 08:40:05 2021 -0700

    ice: mark PTYPE 2 as reserved

    [ Upstream commit 0c526d440f76676733cb470b454db9d5507a3a50 ]

    The entry for PTYPE 2 in the ice_ptype_lkup table incorrectly states
    that this is an L2 packet with no payload. According to the datasheet,
    this PTYPE is actually unused and reserved.

    Fix the lookup entry to indicate this is an unused entry that is
```

```
      reserved.

      Signed-off-by: Jacob Keller <jacob.e.keller@intel.com>
      Tested-by: Tony Brelinski <tonyx.brelinski@intel.com>
      Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 008ec090f0d1f3eb8d462723d9be439065f31b4d
  Author: Jacob Keller <jacob.e.keller@intel.com>
  Date:   Thu May 6 08:40:04 2021 -0700

      ice: fix incorrect payload indicator on PTYPE

      [ Upstream commit 638a0c8c8861cb8a3b54203e632ea5dcc23d8ca5 ]

      The entry for PTYPE 90 indicates that the payload is layer 3. This does
      not match the specification in the datasheet which indicates the packet
      is a MAC, IPv6, UDP packet, with a payload in layer 4.

      Fix the lookup table to match the data sheet.

      Signed-off-by: Jacob Keller <jacob.e.keller@intel.com>
      Tested-by: Tony Brelinski <tonyx.brelinski@intel.com>
      Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit eb7b40d9d3785f7a131fb0b1f89bb6efa46c1833
  Author: Pavel Skripkin <paskripkin@gmail.com>
  Date:   Fri Apr 30 21:50:46 2021 +0300

      ext4: fix memory leak in ext4_fill_super

      [ Upstream commit 618f003199c6188e01472b03cdbba227f1dc5f24 ]

      static int kthread(void *_create) will return -ENOMEM
      or -EINTR in case of internal failure or
      kthread_stop() call happens before threadfn call.

      To prevent fancy error checking and make code
      more straightforward we moved all cleanup code out
      of kmmpd threadfn.

      Also, dropped struct mmpd_data at all. Now struct super_block
      is a threadfn data and struct buffer_head embedded into
      struct ext4_sb_info.

      Reported-by: syzbot+d9e482e303930fa4f6ff@syzkaller.appspotmail.com
      Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
      Link: https://lore.kernel.org/r/20210430185046.15742-1-paskripkin@gmail.com
      Signed-off-by: Theodore Ts'o <tytso@mit.edu>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4d39b8f57b7dd0f2713df849f5126c165089af8f
  Author: Daniel Borkmann <daniel@iogearbox.net>
  Date:   Wed Jun 16 11:25:11 2021 +0200

      bpf: Fix up register-based shifts in interpreter to silence KUBSAN

      [ Upstream commit 28131e9d933339a92f78e7ab6429f4aaaa07061c ]

      syzbot reported a shift-out-of-bounds that KUBSAN observed in the
      interpreter:

        [...]
        UBSAN: shift-out-of-bounds in kernel/bpf/core.c:1420:2
```

```
       shift exponent 255 is too large for 64-bit type 'long long unsigned int'
       CPU: 1 PID: 11097 Comm: syz-executor.4 Not tainted 5.12.0-rc2-syzkaller #0
       Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
 01/01/2011
       Call Trace:
        __dump_stack lib/dump_stack.c:79 [inline]
        dump_stack+0x141/0x1d7 lib/dump_stack.c:120
        ubsan_epilogue+0xb/0x5a lib/ubsan.c:148
        __ubsan_handle_shift_out_of_bounds.cold+0xb1/0x181 lib/ubsan.c:327
        ___bpf_prog_run.cold+0x19/0x56c kernel/bpf/core.c:1420
        __bpf_prog_run32+0x8f/0xd0 kernel/bpf/core.c:1735
        bpf_dispatcher_nop_func include/linux/bpf.h:644 [inline]
        bpf_prog_run_pin_on_cpu include/linux/filter.h:624 [inline]
        bpf_prog_run_clear_cb include/linux/filter.h:755 [inline]
        run_filter+0x1a1/0x470 net/packet/af_packet.c:2031
        packet_rcv+0x313/0x13e0 net/packet/af_packet.c:2104
        dev_queue_xmit_nit+0x7c2/0xa90 net/core/dev.c:2387
        xmit_one net/core/dev.c:3588 [inline]
        dev_hard_start_xmit+0xad/0x920 net/core/dev.c:3609
        __dev_queue_xmit+0x2121/0x2e00 net/core/dev.c:4182
        __bpf_tx_skb net/core/filter.c:2116 [inline]
        __bpf_redirect_no_mac net/core/filter.c:2141 [inline]
        __bpf_redirect+0x548/0xc80 net/core/filter.c:2164
        ____bpf_clone_redirect net/core/filter.c:2448 [inline]
        bpf_clone_redirect+0x2ae/0x420 net/core/filter.c:2420
        ___bpf_prog_run+0x34e1/0x77d0 kernel/bpf/core.c:1523
        __bpf_prog_run512+0x99/0xe0 kernel/bpf/core.c:1737
        bpf_dispatcher_nop_func include/linux/bpf.h:644 [inline]
        bpf_test_run+0x3ed/0xc50 net/bpf/test_run.c:50
        bpf_prog_test_run_skb+0xabc/0x1c50 net/bpf/test_run.c:582
        bpf_prog_test_run kernel/bpf/syscall.c:3127 [inline]
        __do_sys_bpf+0x1ea9/0x4f00 kernel/bpf/syscall.c:4406
        do_syscall_64+0x2d/0x70 arch/x86/entry/common.c:46
        entry_SYSCALL_64_after_hwframe+0x44/0xae
       [...]
```

Generally speaking, KUBSAN reports from the kernel should be fixed. However, in case of BPF, this particular report caused concerns since the large shift is not wrong from BPF point of view, just undefined. In the verifier, K-based shifts that are >= {64,32} (depending on the bitwidth of the instruction) are already rejected. The register-based cases were not given their content might not be known at verification time. Ideas such as verifier instruction rewrite with an additional AND instruction for the source register were brought up, but regularly rejected due to the additional runtime overhead they incur.

As Edward Cree rightly put it:

> Shifts by more than insn bitness are legal in the BPF ISA; they are implementation-defined behaviour [of the underlying architecture], rather than UB, and have been made legal for performance reasons. Each of the JIT backends compiles the BPF shift operations to machine instructions which produce implementation-defined results in such a case; the resulting contents of the register may be arbitrary but program behaviour as a whole remains defined.
>
> Guard checks in the fast path (i.e. affecting JITted code) will thus not be accepted.
>
> The case of division by zero is not truly analogous here, as division instructions on many of the JIT-targeted architectures will raise a machine exception / fault on division by zero, whereas (to the best of my knowledge) none will do so on an out-of-bounds shift.

Given the KUBSAN report only affects the BPF interpreter, but not JITs,

```
    one solution is to add the ANDs with 63 or 31 into ___bpf_prog_run().
    That would make the shifts defined, and thus shuts up KUBSAN, and the
    compiler would optimize out the AND on any CPU that interprets the shift
    amounts modulo the width anyway (e.g., confirmed from disassembly that
    on x86-64 and arm64 the generated interpreter code is the same before
    and after this fix).

    The BPF interpreter is slow path, and most likely compiled out anyway
    as distros select BPF_JIT_ALWAYS_ON to avoid speculative execution of
    BPF instructions by the interpreter. Given the main argument was to
    avoid sacrificing performance, the fact that the AND is optimized away
    from compiler for mainstream archs helps as well as a solution moving
    forward. Also add a comment on LSH/RSH/ARSH translation for JIT authors
    to provide guidance when they see the ___bpf_prog_run() interpreter
    code and use it as a model for a new JIT backend.

    Reported-by: syzbot+bed360704c521841c85d@syzkaller.appspotmail.com
    Reported-by: Kurt Manucredo <fuzzybritches0@gmail.com>
    Signed-off-by: Eric Biggers <ebiggers@kernel.org>
    Co-developed-by: Eric Biggers <ebiggers@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Andrii Nakryiko <andrii@kernel.org>
    Tested-by: syzbot+bed360704c521841c85d@syzkaller.appspotmail.com
    Cc: Edward Cree <ecree.xilinx@gmail.com>
    Link: https://lore.kernel.org/bpf/0000000000008f912605bd30d5d7@google.com
    Link: https://lore.kernel.org/bpf/bac16d8d-c174-bdc4-91bd-bfa62b410190@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2cb85397c30f5152ad7d8f1414f9635a0e23bc2f
Author: George McCollister <george.mccollister@gmail.com>
Date:   Tue Jun 15 12:50:37 2021 -0500

    net: hsr: don't check sequence number if tag removal is offloaded

    [ Upstream commit c2ae34a7deaff463ecafb7db627b77faaca8e159 ]

    Don't check the sequence number when deciding when to update time_in in
    the node table if tag removal is offloaded since the sequence number is
    part of the tag. This fixes a problem where the times in the node table
    wouldn't update when 0 appeared to be before or equal to seq_out when
    tag removal was offloaded.

    Signed-off-by: George McCollister <george.mccollister@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3382ec34a4d56c02a2c146e9c6c73bfe8d6e2604
Author: Amber Lin <Amber.Lin@amd.com>
Date:   Mon Jun 7 14:46:21 2021 -0400

    drm/amdkfd: Fix circular lock in nocpsch path

    [ Upstream commit a7b2451d31cfa2e8aeccf3b35612ce33f02371fc ]

    Calling free_mqd inside of destroy_queue_nocpsch_locked can cause a
    circular lock. destroy_queue_nocpsch_locked is called under a DQM lock,
    which is taken in MMU notifiers, potentially in FS reclaim context.
    Taking another lock, which is BO reservation lock from free_mqd, while
    causing an FS reclaim inside the DQM lock creates a problematic circular
    lock dependency. Therefore move free_mqd out of
    destroy_queue_nocpsch_locked and call it after unlocking DQM.

    Signed-off-by: Amber Lin <Amber.Lin@amd.com>
    Reviewed-by: Felix Kuehling <Felix.Kuehling@amd.com>
```

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

commit 161c4c2a19f494cb2aceeaa315b442abab649f7a
Author: Jonathan Kim <jonathan.kim@amd.com>
Date:   Fri Jun 11 13:36:34 2021 -0400

    drm/amdkfd: fix circular locking on get_wave_state

    [ Upstream commit 63f6e01237257e7226efc5087f3f0b525d320f54 ]

    get_wave_state acquires the mmap_lock on copy_to_user but so do
    mmu_notifiers.  mmu_notifiers allows dqm locking so do get_wave_state
    outside the dqm_lock to prevent circular locking.

    v2: squash in unused variable removal.

    Signed-off-by: Jonathan Kim <jonathan.kim@amd.com>
    Reviewed-by: Felix Kuehling <felix.kuehling@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 118e6aea7ae7c92ac85ec37dc4d0a6111c0e2a37
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 11:05:14 2021 +0800

    cw1200: add missing MODULE_DEVICE_TABLE

    [ Upstream commit dd778f89225cd258e8f0fed2b7256124982c8bb5 ]

    This patch adds missing MODULE_DEVICE_TABLE definition which generates
    correct modalias for automatic loading of this driver when it is built
    as an external module.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/1620788714-14300-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 40af3960a15339e8bbd3be50c3bc7b35e1a0b6ea
Author: Lee Gibson <leegib@gmail.com>
Date:   Wed Apr 28 12:55:08 2021 +0100

    wl1251: Fix possible buffer overflow in wl1251_cmd_scan

    [ Upstream commit d10a87a3535cce2b890897914f5d0d83df669c63 ]

    Function wl1251_cmd_scan calls memcpy without checking the length.
    Harden by checking the length is within the maximum allowed size.

    Signed-off-by: Lee Gibson <leegib@gmail.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210428115508.25624-1-leegib@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 36409fe65122a51ea0a9f14e6a9b4206eba707b6
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Mon Jun 14 16:44:41 2021 +0300

    net: phy: nxp-c45-tja11xx: enable MDIO write access to the master/slave registers

    [ Upstream commit 0b5f0f29b118910c89fe249cdfbc11b400a86a18 ]

    The SJA1110 switch integrates TJA1103 PHYs, but in SJA1110 switch rev B
```

```
    silicon, there is a bug in that the registers for selecting the 100base-T1
    autoneg master/slave roles are not writable.

    To enable write access to the master/slave registers, these additional
    PHY writes are necessary during initialization.

    The issue has been corrected in later SJA1110 silicon versions and is
    not present in the standalone PHY variants, but applying the workaround
    unconditionally in the driver should not do any harm.

    Suggested-by: Radu Pirea (NXP OSS) <radu-nicolae.pirea@oss.nxp.com>
    Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Reviewed-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84bb21d9886c70b466b919c454c6bbd52e6b5e41
Author: Tony Lindgren <tony@atomide.com>
Date:   Thu Jun 3 09:28:14 2021 +0300

    wlcore/wl12xx: Fix wl12xx get_mac error if device is in ELP

    [ Upstream commit 11ef6bc846dcdce838f0b00c5f6a562c57e5d43b ]

    At least on wl12xx, reading the MAC after boot can fail with a warning
    at drivers/net/wireless/ti/wlcore/sdio.c:78 wl12xx_sdio_raw_read.
    The failed call comes from wl12xx_get_mac() that wlcore_nvs_cb() calls
    after request_firmware_work_func().

    After the error, no wireless interface is created. Reloading the wl12xx
    module makes the interface work.

    Turns out the wlan controller can be in a low-power ELP state after the
    boot from the bootloader or kexec, and needs to be woken up first.

    Let's wake the hardware and add a sleep after that similar to
    wl12xx_pre_boot() is already doing.

    Note that a similar issue could exist for wl18xx, but I have not seen it
    so far. And a search for wl18xx_get_mac and wl12xx_sdio_raw_read did not
    produce similar errors.

    Cc: Carl Philipp Klemm <philipp@uvos.xyz>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210603062814.19464-1-tony@atomide.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce6dc850612a6cb8d2980222804cf8839dc7758a
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Sun Jun 6 16:13:16 2021 -0400

    dm writecache: commit just one block, not a full page

    [ Upstream commit 991bd8d7bc78966b4dc427b53a144f276bffcd52 ]

    Some architectures have pages larger than 4k and committing a full
    page causes needless overhead.

    Fix this by writing a single block when committing the superblock.

    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 7cc35bd4355b5953c6faa9c067ebbc8ef9c84aa7
Author: Steffen Klassert <steffen.klassert@secunet.com>
Date:   Mon Jun 7 15:21:49 2021 +0200

    xfrm: Fix error reporting in xfrm_state_construct.

    [ Upstream commit 6fd06963fa74197103cdbb4b494763127b3f2f34 ]

    When memory allocation for XFRMA_ENCAP or XFRMA_COADDR fails,
    the error will not be reported because the -ENOMEM assignment
    to the err variable is overwritten before. Fix this by moving
    these two in front of the function so that memory allocation
    failures will be reported.

    Reported-by: Tobias Brunner <tobias@strongswan.org>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 62adf2af9f09e235dfb5621a8858b7d22f105a41
Author: Lijun Pan <lijunp213@gmail.com>
Date:   Fri Jun 11 10:43:39 2021 -0500

    ibmvnic: fix kernel build warnings in build_hdr_descs_arr

    [ Upstream commit 73214a690c50a134bd364e1a4430e0e7ac81a8d8 ]

    Fix the following kernel build warnings:
    drivers/net/ethernet/ibm/ibmvnic.c:1516: warning: Function parameter or member 'skb' not
described in 'build_hdr_descs_arr'
    drivers/net/ethernet/ibm/ibmvnic.c:1516: warning: Function parameter or member
'indir_arr' not described in 'build_hdr_descs_arr'
    drivers/net/ethernet/ibm/ibmvnic.c:1516: warning: Excess function parameter 'txbuff'
description in 'build_hdr_descs_arr'

    Signed-off-by: Lijun Pan <lijunp213@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20102b0f678493b95ea1f142246fff726ac7cc4c
Author: Mark Yacoub <markyacoub@chromium.org>
Date:   Fri Jun 4 13:01:07 2021 -0400

    drm/amd/display: Verify Gamma & Degamma LUT sizes in amdgpu_dm_atomic_check

    [ Upstream commit 03fc4cf45d30533d54f0f4ebc02aacfa12f52ce2 ]

    For each CRTC state, check the size of Gamma and Degamma LUTs  so
    unexpected and larger sizes wouldn't slip through.

    TEST: IGT:kms_color::pipe-invalid-gamma-lut-sizes

    v2: fix assignments in if clauses, Mark's email.

    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Mark Yacoub <markyacoub@chromium.org>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66fc2862a252f7356e7aa5eeb44519ee005fd066
Author: Heiner Kallweit <hkallweit1@gmail.com>
Date:   Thu Jun 10 22:56:59 2021 +0200

    r8169: avoid link-up interrupt issue on RTL8106e if user enables ASPM

    [ Upstream commit 1ee8856de82faec9bc8bd0f2308a7f27e30ba207 ]
```

```
    It has been reported that on RTL8106e the link-up interrupt may be
    significantly delayed if the user enables ASPM L1. Per default ASPM
    is disabled. The change leaves L1 enabled on the PCIe link (thus still
    allowing to reach higher package power saving states), but the
    NIC won't actively trigger it.

    Reported-by: Koba Ko <koba.ko@canonical.com>
    Tested-by: Koba Ko <koba.ko@canonical.com>
    Signed-off-by: Heiner Kallweit <hkallweit1@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4b58c2891d09b07225cb8c4199de62dbb6a1b441
Author: Minchan Kim <minchan@kernel.org>
Date:   Wed Jun 9 09:37:17 2021 -0700

    selinux: use __GFP_NOWARN with GFP_NOWAIT in the AVC

    [ Upstream commit 648f2c6100cfa18e7dfe43bc0b9c3b73560d623c ]

    In the field, we have seen lots of allocation failure from the call
    path below.

    06-03 13:29:12.999 1010315 31557 31557 W Binder  : 31542_2: page allocation failure:
order:0, mode:0x800(GFP_NOWAIT), nodemask=(null),cpuset=background,mems_allowed=0
    ...
    ...
    06-03 13:29:12.999 1010315 31557 31557 W Call trace:
    06-03 13:29:12.999 1010315 31557 31557 W         : dump_backtrace.cfi_jt+0x0/0x8
    06-03 13:29:12.999 1010315 31557 31557 W         : dump_stack+0xc8/0x14c
    06-03 13:29:12.999 1010315 31557 31557 W         : warn_alloc+0x158/0x1c8
    06-03 13:29:12.999 1010315 31557 31557 W         : __alloc_pages_slowpath+0x9d8/0xb80
    06-03 13:29:12.999 1010315 31557 31557 W         : __alloc_pages_nodemask+0x1c4/0x430
    06-03 13:29:12.999 1010315 31557 31557 W         : allocate_slab+0xb4/0x390
    06-03 13:29:12.999 1010315 31557 31557 W         : ___slab_alloc+0x12c/0x3a4
    06-03 13:29:12.999 1010315 31557 31557 W         : kmem_cache_alloc+0x358/0x5e4
    06-03 13:29:12.999 1010315 31557 31557 W         : avc_alloc_node+0x30/0x184
    06-03 13:29:12.999 1010315 31557 31557 W         : avc_update_node+0x54/0x4f0
    06-03 13:29:12.999 1010315 31557 31557 W         : avc_has_extended_perms+0x1a4/0x460
    06-03 13:29:12.999 1010315 31557 31557 W         : selinux_file_ioctl+0x320/0x3d0
    06-03 13:29:12.999 1010315 31557 31557 W         : __arm64_sys_ioctl+0xec/0x1fc
    06-03 13:29:12.999 1010315 31557 31557 W         : el0_svc_common+0xc0/0x24c
    06-03 13:29:12.999 1010315 31557 31557 W         : el0_svc+0x28/0x88
    06-03 13:29:12.999 1010315 31557 31557 W         : el0_sync_handler+0x8c/0xf0
    06-03 13:29:12.999 1010315 31557 31557 W         : el0_sync+0x1a4/0x1c0
    ..
    ..
    06-03 13:29:12.999 1010315 31557 31557 W SLUB    : Unable to allocate memory on node -1,
gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
    06-03 13:29:12.999 1010315 31557 31557 W cache   : avc_node, object size: 72, buffer
size: 80, default order: 0, min order: 0
    06-03 13:29:12.999 1010315 31557 31557 W node 0  : slabs: 57, objs: 2907, free: 0
    06-03 13:29:12.999 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
    06-03 13:29:12.999 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
size: 80, default order: 0, min order: 0
    06-03 13:29:12.999 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
    06-03 13:29:12.999 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
    06-03 13:29:12.999 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
size: 80, default order: 0, min order: 0
    06-03 13:29:12.999 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
    06-03 13:29:12.999 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
```

```
      06-03 13:29:12.999 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
  size: 80, default order: 0, min order: 0
      06-03 13:29:12.999 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
      06-03 13:29:13.000 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
  gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
      06-03 13:29:13.000 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
  size: 80, default order: 0, min order: 0
      06-03 13:29:13.000 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
      06-03 13:29:13.000 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
  gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
      06-03 13:29:13.000 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
  size: 80, default order: 0, min order: 0
      06-03 13:29:13.000 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
      06-03 13:29:13.000 1010161 10686 10686 W SLUB    : Unable to allocate memory on node -1,
  gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
      06-03 13:29:13.000 1010161 10686 10686 W cache   : avc_node, object size: 72, buffer
  size: 80, default order: 0, min order: 0
      06-03 13:29:13.000 1010161 10686 10686 W node 0  : slabs: 57, objs: 2907, free: 0
      06-03 13:29:13.000 10230 30892 30892 W SLUB    : Unable to allocate memory on node -1,
  gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
      06-03 13:29:13.000 10230 30892 30892 W cache   : avc_node, object size: 72, buffer size:
  80, default order: 0, min order: 0
      06-03 13:29:13.000 10230 30892 30892 W node 0  : slabs: 57, objs: 2907, free: 0
      06-03 13:29:13.000 10230 30892 30892 W SLUB    : Unable to allocate memory on node -1,
  gfp=0x900(GFP_NOWAIT|__GFP_ZERO)
      06-03 13:29:13.000 10230 30892 30892 W cache   : avc_node, object size: 72, buffer size:
  80, default order: 0, min order: 0

      Based on [1], selinux is tolerate for failure of memory allocation.
      Then, use __GFP_NOWARN together.

      [1] 476accbe2f6e ("selinux: use GFP_NOWAIT in the AVC kmem_caches")

      Signed-off-by: Minchan Kim <minchan@kernel.org>
      [PM: subj fix, line wraps, normalized commit refs]
      Signed-off-by: Paul Moore <paul@paul-moore.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit dd18ad96a9ea267152eec034e24efbdc7eb7df2f
  Author: Yang Yingliang <yangyingliang@huawei.com>
  Date:   Thu Jun 10 17:17:12 2021 +0800

      net: mido: mdio-mux-bcm-iproc: Use devm_platform_get_and_ioremap_resource()

      [ Upstream commit 8a55a73433e763c8aec4a3e8df5c28c821fc44b9 ]

      Use devm_platform_get_and_ioremap_resource() to simplify
      code and avoid a null-ptr-deref by checking 'res' in it.

      Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 7ca7d4404b959a4642908bbb227de480aa57ba29
  Author: Yang Yingliang <yangyingliang@huawei.com>
  Date:   Thu Jun 10 16:02:43 2021 +0800

      fjes: check return value after calling platform_get_resource()

      [ Upstream commit f18c11812c949553d2b2481ecaa274dd51bed1e7 ]

      It will cause null-ptr-deref if platform_get_resource() returns NULL,
      we need check the return value.

      Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
```

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

commit 61e364fafec1d31b68c61653e81c577f471ee780
Author: Nirmoy Das <nirmoy.das@amd.com>
Date:   Tue Jun 8 13:23:44 2021 +0200

    drm/amdkfd: use allowed domain for vmbo validation

    [ Upstream commit bc05716d4fdd065013633602c5960a2bf1511b9c ]

    Fixes handling when page tables are in system memory.

    v3: remove struct amdgpu_vm_parser.
    v2: remove unwanted variable.
        change amdgpu_amdkfd_validate instead of amdgpu_amdkfd_bo_validate.

    Signed-off-by: Nirmoy Das <nirmoy.das@amd.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Reviewed-by: Felix Kuehling <Felix.Kuehling@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b5c55df50679c7fde8eb424afe2925eaf38e7aa
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Jun 9 21:25:15 2021 +0800

    net: sgi: ioc3-eth: check return value after calling platform_get_resource()

    [ Upstream commit db8f7be1e1d64fbf113a456ef94534fbf5e9a9af ]

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 42b8412dc12309c3050b38818e478e24e186c647
Author: Amit Cohen <amcohen@nvidia.com>
Date:   Tue Jun 8 15:44:09 2021 +0300

    selftests: Clean forgotten resources as part of cleanup()

    [ Upstream commit e67dfb8d15deb33c425d0b0ee22f2e5eef54c162 ]

    Several tests do not set some ports down as part of their cleanup(),
    resulting in IPv6 link-local addresses and associated routes not being
    deleted.

    These leaks were found using a BPF tool that monitors ASIC resources.

    Solve this by setting the ports down at the end of the tests.

    Signed-off-by: Amit Cohen <amcohen@nvidia.com>
    Reviewed-by: Petr Machata <petrm@nvidia.com>
    Signed-off-by: Ido Schimmel <idosch@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32438665ba6634ff7b19bad3305934096cc0acee
Author: Joakim Zhang <qiangqing.zhang@nxp.com>
Date:   Tue Jun 8 11:15:35 2021 +0800

    net: phy: realtek: add delay to fix RXC generation issue
```

```
    [ Upstream commit 6813cc8cfdaf401476e1a007cec8ae338cefa573 ]

    PHY will delay about 11.5ms to generate RXC clock when switching from
    power down to normal operation. Read/write registers would also cause RXC
    become unstable and stop for a while during this process. Realtek engineer
    suggests 15ms or more delay can workaround this issue.

    Signed-off-by: Joakim Zhang <qiangqing.zhang@nxp.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3df8a157824b91b4a1ed793cd942d6df038a7b72
Author: Aric Cyr <aric.cyr@amd.com>
Date:   Tue May 25 15:33:51 2021 -0400

    drm/amd/display: Fix crash during MPO + ODM combine mode recalculation

    [ Upstream commit 665f28507a2a3d8d72ed9afa9a2b9b17fd43add1 ]

    [Why]
    When calculating recout width for an MPO plane on a mode that's using
    ODM combine, driver can calculate a negative value, resulting in a
    crash.

    [How]
    For negative widths, use zero such that validation will prune the
    configuration correctly and disallow MPO.

    Signed-off-by: Aric Cyr <aric.cyr@amd.com>
    Reviewed-by: Krunoslav Kovac <Krunoslav.Kovac@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit caef3b95ce82616b168c9ab287bce4bb6a7aa2f1
Author: Wesley Chalmers <Wesley.Chalmers@amd.com>
Date:   Wed Apr 28 18:38:54 2021 -0400

    drm/amd/display: Fix off-by-one error in DML

    [ Upstream commit e4e3678260e9734f6f41b4325aac0b171833a618 ]

    [WHY]
    For DCN30 and later, there is no data in DML arrays indexed by state at
    index num_states.

    Signed-off-by: Wesley Chalmers <Wesley.Chalmers@amd.com>
    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f88bc77bbc75b25df55d1da744812b018f6830a
Author: Wesley Chalmers <Wesley.Chalmers@amd.com>
Date:   Thu May 20 12:12:48 2021 -0400

    drm/amd/display: Set DISPCLK_MAX_ERRDET_CYCLES to 7

    [ Upstream commit 3577e1678772ce3ede92af3a75b44a4b76f9b4ad ]

    [WHY]
    DISPCLK_MAX_ERRDET_CYCLES must be 7 to prevent connection loss when
```

```
    changing DENTIST_DISPCLK_WDIVIDER from 126 to 127 and back.

    Signed-off-by: Wesley Chalmers <Wesley.Chalmers@amd.com>
    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db1ca5441e33273d37a2e3b9d90e7dc871447e83
Author: Vladimir Stempen <vladimir.stempen@amd.com>
Date:   Wed May 19 13:55:46 2021 -0400

    drm/amd/display: Release MST resources on switch from MST to SST

    [ Upstream commit 3f8518b60c10aa96f3efa38a967a0b4eb9211ac0 ]

    [why]
    When OS overrides training link training parameters
    for MST device to SST mode, MST resources are not
    released and leak of the resource may result crash and
    incorrect MST discovery during following hot plugs.

    [how]
    Retaining sink object to be reused by SST link and
    releasing MST  resources.

    Signed-off-by: Vladimir Stempen <vladimir.stempen@amd.com>
    Reviewed-by: Wenjing Liu <Wenjing.Liu@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71e0e437e97b3312b6377d6079667586f51afaf8
Author: Roman Li <roman.li@amd.com>
Date:   Fri May 21 10:20:25 2021 -0400

    drm/amd/display: Update scaling settings on modeset

    [ Upstream commit c521fc316d12fb9ea7b7680e301d673bceda922e ]

    [Why]
    We update scaling settings when scaling mode has been changed.
    However when changing mode from native resolution the scaling mode previously
    set gets ignored.

    [How]
    Perform scaling settings update on modeset.

    Signed-off-by: Roman Li <roman.li@amd.com>
    Reviewed-by: Nicholas Kazlauskas <Nicholas.Kazlauskas@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc9e10ea5bdd56bdb0631031adca5846ed6103d7
Author: Ilya Bakoulin <Ilya.Bakoulin@amd.com>
Date:   Tue May 18 17:55:17 2021 -0400

    drm/amd/display: Revert "Fix clock table filling logic"

    [ Upstream commit ae88357c7966ec2a52cb1e70fd42f74a40c9dfcb ]
```

```
    [Why]
    This change was found to break some high-refresh modes. Reverting
    to unblock mainline.

    Signed-off-by: Ilya Bakoulin <Ilya.Bakoulin@amd.com>
    Reviewed-by: Sung Lee <Sung.Lee@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef485944e718088f99fdc2a791efd4c88dd0ac9a
Author: Nikola Cornij <nikola.cornij@amd.com>
Date:   Tue May 18 12:33:12 2021 -0400

    drm/amd/display: Fix DCN 3.01 DSCCLK validation

    [ Upstream commit 346cf627fb27c0fea63a041cedbaa4f31784e504 ]

    [why]
    DSCCLK validation is not necessary because DSCCLK is derrived from
    DISPCLK, therefore if DISPCLK validation passes, DSCCLK is valid, too.
    Doing DSCLK validation in addition to DISPCLK leads to modes being
    wrongly rejected when DSCCLK was incorrectly set outside of DML.

    [how]
    Remove DSCCLK validation because it's implicitly validated under DISPCLK

    Signed-off-by: Nikola Cornij <nikola.cornij@amd.com>
    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 18aa107512275a7bc467fb3229193d866653a772
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Mon Jun 7 23:02:59 2021 +0800

    net: moxa: Use devm_platform_get_and_ioremap_resource()

    [ Upstream commit 35cba15a504bf4f585bb9d78f47b22b28a1a06b2 ]

    Use devm_platform_get_and_ioremap_resource() to simplify
    code and avoid a null-ptr-deref by checking 'res' in it.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c01540159ead3be3bbf0edbe23db0de04453b481
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Mon Jun 7 22:55:21 2021 +0800

    net: micrel: check return value after calling platform_get_resource()

    [ Upstream commit 20f1932e2282c58cb5ac59517585206cf5b385ae ]

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 373db5a9c8c4a934e89f077f1077bb74c37c1bd6
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Mon Jun 7 22:36:02 2021 +0800

    net: mvpp2: check return value after calling platform_get_resource()

    [ Upstream commit 0bb51a3a385790a4be20085494cf78f70dadf646 ]

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd9575caf7bdf5631328c082bba0c734cf9f87cf
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Mon Jun 7 21:38:37 2021 +0800

    net: bcmgenet: check return value after calling platform_get_resource()

    [ Upstream commit 74325bf0104573c6dfce42837139aeef3f34be76 ]

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 43630cbcd7de3345d32fbe5bac1063e484a95286
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Sat Jun 5 10:31:48 2021 +0800

    net: mscc: ocelot: check return value after calling platform_get_resource()

    [ Upstream commit f1fe19c2cb3fdc92a614cf330ced1613f8f1a681 ]

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Reviewed-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8f03933f5a7a5dfcbaeb113e6fe530b4e3a4763
Author: Xianting Tian <xianting.tian@linux.alibaba.com>
Date:   Sat Jun 5 11:31:00 2021 -0400

    virtio_net: Remove BUG() to avoid machine dead

    [ Upstream commit 85eb1389458d134bdb75dad502cc026c3753a619 ]

    We should not directly BUG() when there is hdr error, it is
    better to output a print when such error happens. Currently,
    the caller of xmit_skb() already did it.

    Signed-off-by: Xianting Tian <xianting.tian@linux.alibaba.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47cafc568c31e3d3179cc51cd2ca4a696adaddd3
```

```
Author: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
Date:   Wed Mar 31 14:17:06 2021 -0700

    ice: fix clang warning regarding deadcode.DeadStores

    [ Upstream commit 7e94090ae13e1ae5fe8bd3a9cd08136260bb7039 ]

    clang generates deadcode.DeadStores warnings when a variable
    is used to read a value, but then that value isn't used later
    in the code. Fix this warning.

    Signed-off-by: Paul M Stillwell Jr <paul.m.stillwell.jr@intel.com>
    Tested-by: Tony Brelinski <tonyx.brelinski@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63efeeb1274d44e3c490be3b68cc74de5d63ac79
Author: Liwei Song <liwei.song@windriver.com>
Date:   Mon Apr 19 17:31:06 2021 +0800

    ice: set the value of global config lock timeout longer

    [ Upstream commit fb3612840d4f587a0af9511a11d7989d1fa48206 ]

    It may need hold Global Config Lock a longer time when download DDP
    package file, extend the timeout value to 5000ms to ensure that
    download can be finished before other AQ command got time to run,
    this will fix the issue below when probe the device, 5000ms is a test
    value that work with both Backplane and BreakoutCable NVM image:

    ice 0000:f4:00.0: VSI 12 failed lan queue config, error ICE_ERR_CFG
    ice 0000:f4:00.0: Failed to delete VSI 12 in FW - error: ICE_ERR_AQ_TIMEOUT
    ice 0000:f4:00.0: probe failed due to setup PF switch: -12
    ice: probe of 0000:f4:00.0 failed with error -12

    Signed-off-by: Liwei Song <liwei.song@windriver.com>
    Tested-by: Tony Brelinski <tonyx.brelinski@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22f5c64be9c790968d4f87a0c801fe8124fc1141
Author: Radim Pavlik <radim.pavlik@tbs-biometrics.com>
Date:   Tue Jun 1 10:48:18 2021 +0000

    pinctrl: mcp23s08: fix race condition in irq handler

    [ Upstream commit 897120d41e7afd9da435cb00041a142aeeb53c07 ]

    Checking value of MCP_INTF in mcp23s08_irq suggests that the handler may be
    called even when there is no interrupt pending.

    But the actual interrupt could happened between reading MCP_INTF and MCP_GPIO.
    In this situation we got nothing from MCP_INTF, but the event gets acknowledged
    on the expander by reading MCP_GPIO. This leads to losing events.

    Fix the problem by not reading any register until we see something in MCP_INTF.

    The error was reproduced and fix tested on MCP23017.

    Signed-off-by: Radim Pavlik <radim.pavlik@tbs-biometrics.com>
    Link:
https://lore.kernel.org/r/AM7PR06MB6769E1183F68DEBB252F665ABA3E9@AM7PR06MB6769.eurprd06.prod
.outlook.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 88e1357e4d0504c2e8159daeb5b205e8ca52120f
Author: Horatiu Vultur <horatiu.vultur@microchip.com>
Date:   Fri Jun 4 12:37:47 2021 +0200

    net: bridge: mrp: Update ring transitions.

    [ Upstream commit fcb34635854a5a5814227628867ea914a9805384 ]

    According to the standard IEC 62439-2, the number of transitions needs
    to be counted for each transition 'between' ring state open and ring
    state closed and not from open state to closed state.

    Therefore fix this for both ring and interconnect ring.

    Signed-off-by: Horatiu Vultur <horatiu.vultur@microchip.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03de5ab1dc405d53dca149427fa4eaaaea7ac783
Author: Damien Le Moal <damien.lemoal@wdc.com>
Date:   Wed May 26 06:24:53 2021 +0900

    block: introduce BIO_ZONE_WRITE_LOCKED bio flag

    [ Upstream commit 9ffbbb435d8f566a0924ce4b5dc7fc1bceb6dbf8 ]

    Introduce the BIO flag BIO_ZONE_WRITE_LOCKED to indicate that a BIO owns
    the write lock of the zone it is targeting. This is the counterpart of
    the struct request flag RQF_ZONE_WRITE_LOCKED.

    This new BIO flag is reserved for now for zone write locking control
    for device mapper targets exposing a zoned block device. Since in this
    case, the lock flag must not be propagated to the struct request that
    will be used to process the BIO, a BIO private flag is used rather than
    changing the RQF_ZONE_WRITE_LOCKED request flag into a common REQ_XXX
    flag that could be used for both BIO and request. This avoids conflicts
    down the stack with the block IO scheduler zone write locking
    (in mq-deadline).

    Signed-off-by: Damien Le Moal <damien.lemoal@wdc.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Reviewed-by: Chaitanya Kulkarni <chaitanya.kulkarni@wdc.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Acked-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4068cb7e11131009527f6f134df098b1357d8748
Author: Damien Le Moal <damien.lemoal@wdc.com>
Date:   Wed May 26 06:24:54 2021 +0900

    dm: Fix dm_accept_partial_bio() relative to zone management commands

    [ Upstream commit 6842d264aa5205da338b6dcc6acfa2a6732558f1 ]

    Fix dm_accept_partial_bio() to actually check that zone management
    commands are not passed as explained in the function documentation
    comment. Also, since a zone append operation cannot be split, add
    REQ_OP_ZONE_APPEND as a forbidden command.

    White lines are added around the group of BUG_ON() calls to make the
    code more legible.
```

```
    Signed-off-by: Damien Le Moal <damien.lemoal@wdc.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a8a10887b6537e1c6b534a310888fb1213a0708
Author: Mikulas Patocka <mpatocka@redhat.com>
Date:   Wed May 26 10:29:45 2021 -0400

    dm writecache: don't split bios when overwriting contiguous cache content

    [ Upstream commit ee50cc19d80e9b9a8283d1fb517a778faf2f6899 ]

    If dm-writecache overwrites existing cached data, it splits the
    incoming bio into many block-sized bios. The I/O scheduler does merge
    these bios into one large request but this needless splitting and
    merging causes performance degradation.

    Fix this by avoiding bio splitting if the cache target area that is
    being overwritten is contiguous.

    Signed-off-by: Mikulas Patocka <mpatocka@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c8afdbb62ee7b490b711d62215e1892933e13e1
Author: Joe Thornber <ejt@redhat.com>
Date:   Tue Apr 13 09:03:49 2021 +0100

    dm space maps: don't reset space map allocation cursor when committing

    [ Upstream commit 5faafc77f7de69147d1e818026b9a0cbf036a7b2 ]

    Current commit code resets the place where the search for free blocks
    will begin back to the start of the metadata device.  There are a couple
    of repercussions to this:

    - The first allocation after the commit is likely to take longer than
      normal as it searches for a free block in an area that is likely to
      have very few free blocks (if any).

    - Any free blocks it finds will have been recently freed.  Reusing them
      means we have fewer old copies of the metadata to aid recovery from
      hardware error.

    Fix these issues by leaving the cursor alone, only resetting when the
    search hits the end of the metadata device.

    Signed-off-by: Joe Thornber <ejt@redhat.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c5cc91094e6a4a5363339e0a038e0038e711420f
Author: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
Date:   Tue Jun 1 19:07:49 2021 +0800

    RDMA/cxgb4: Fix missing error code in create_qp()

    [ Upstream commit aeb27bb76ad8197eb47890b1ff470d5faf8ec9a5 ]

    The error code is missing in this code scenario so 0 will be returned. Add
    the error code '-EINVAL' to the return value 'ret'.

    Eliminates the follow smatch warning:

    drivers/infiniband/hw/cxgb4/qp.c:298 create_qp() warn: missing error code 'ret'.
```

```
    Link: https://lore.kernel.org/r/1622545669-20625-1-git-send-email-
jiapeng.chong@linux.alibaba.com
    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Signed-off-by: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 957d6d796e73d2667741a3180741256e60b83d00
Author: Andreas Roeseler <andreas.a.roeseler@gmail.com>
Date:   Thu Jun 3 16:22:11 2021 -0500

    icmp: fix lib conflict with trinity

    [ Upstream commit e32ea44c7ae476f4c90e35ab0a29dc8ff082bc11 ]

    Including <linux/in.h> and <netinet/in.h> in the dependencies breaks
    compilation of trinity due to multiple definitions. <linux/in.h> is only
    used in <linux/icmp.h> to provide the definition of the struct in_addr,
    but this can be substituted out by using the datatype __be32.

    Signed-off-by: Andreas Roeseler <andreas.a.roeseler@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0b8274cda127b1f61f80ea5f7c22b698c49d043
Author: Yuchung Cheng <ycheng@google.com>
Date:   Wed Jun 2 17:51:21 2021 -0700

    net: tcp better handling of reordering then loss cases

    [ Upstream commit a29cb6914681a55667436a9eb7a42e28da8cf387 ]

    This patch aims to improve the situation when reordering and loss are
    ocurring in the same flight of packets.

    Previously the reordering would first induce a spurious recovery, then
    the subsequent ACK may undo the cwnd (based on the timestamps e.g.).
    However the current loss recovery does not proceed to invoke
    RACK to install a reordering timer. If some packets are also lost, this
    may lead to a long RTO-based recovery. An example is
    https://groups.google.com/g/bbr-dev/c/OFHADvJbTEI

    The solution is to after reverting the recovery, always invoke RACK
    to either mount the RACK timer to fast retransmit after the reordering
    window, or restarts the recovery if new loss is identified. Hence
    it is possible the sender may go from Recovery to Disorder/Open to
    Recovery again in one ACK.

    Reported-by: mingkun bian <bianmingkun@gmail.com>
    Signed-off-by: Yuchung Cheng <ycheng@google.com>
    Signed-off-by: Neal Cardwell <ncardwell@google.com>
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b01c274e3a3ab423bcff87f872371715adf22ebd
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue May 18 12:42:47 2021 +0800

    clk: tegra: tegra124-emc: Fix clock imbalance in emc_set_timing()

    [ Upstream commit f13570e7e830ca4fbf4869015af8492b8918445e ]

    After calling clk_prepare_enable(), clk_disable_unprepare() needs
```

```
    be called when prepare_timing_change() failed.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3da33bf6aaffdc48d3fd725fcf4ca19e289083df
Author: Jiansong Chen <Jiansong.Chen@amd.com>
Date:   Thu May 13 10:56:50 2021 +0800

    drm/amdgpu: remove unsafe optimization to drop preamble ib

    [ Upstream commit 7d9c70d23550eb86a1bec1954ccaa8d6ec3a3328 ]

    Take the situation with gfxoff, the optimization may cause
    corrupt CE ram contents. In addition emit_cntxcntl callback
    has similar optimization which firmware can handle properly
    even for power feature.

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

commit 3b2b93a485fb7a970bc8b5daef16f4cf579d172f
Author: Kees Cook <keescook@chromium.org>
Date:   Fri May 28 10:53:54 2021 -0700

    drm/amd/display: Avoid HDCP over-read and corruption

    [ Upstream commit 06888d571b513cbfc0b41949948def6cb81021b2 ]

    Instead of reading the desired 5 bytes of the actual target field,
    the code was reading 8. This could result in a corrupted value if the
    trailing 3 bytes were non-zero, so instead use an appropriately sized
    and zero-initialized bounce buffer, and read only 5 bytes before casting
    to u64.

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

commit 060a3b127d9ed104a0f3e20255926af9ff8370fa
Author: Kevin Wang <kevin1.wang@amd.com>
Date:   Mon May 31 17:29:11 2021 +0800

    drm/amdgpu: fix sdma firmware version error in sriov

    [ Upstream commit 2b8f731849800e3948763ccaff31cceac526789b ]

    Re-adjust the function return order to avoid empty sdma version in the
    sriov environment. (read amdgpu_firmware_info)

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

commit 3e7b75192e1500215726a2ab35f1b578b1c19840
Author: Shiwu Zhang <shiwu.zhang@amd.com>
Date:   Mon May 24 14:30:21 2021 +0800

    drm/amdgpu: fix metadata_size for ubo ioctl queries
```

```
    [ Upstream commit eba98523724be7ad3539f2c975de1527e0c99dd6 ]

    Although the kfd_ioctl_get_dmabuf_info() still fail it will indicate
    the caller right metadat_size useful for the same kfd ioctl next time.

    Signed-off-by: Shiwu Zhang <shiwu.zhang@amd.com>
    Reviewed-by: Nirmoy Das <nirmoy.das@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c77db5a26db6b6399cdefe34ba4bffddeeda244
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sun May 30 18:17:59 2021 +0100

    MIPS: ingenic: Select CPU_SUPPORTS_CPUFREQ && MIPS_EXTERNAL_TIMER

    [ Upstream commit eb3849370ae32b571e1f9a63ba52c61adeaf88f7 ]

    The clock driving the XBurst CPUs in Ingenic SoCs is integer divided
    from the main PLL. As such, it is possible to control the frequency of
    the CPU, either by changing the divider, or by changing the rate of the
    main PLL.

    The XBurst CPUs also lack the CP0 timer; the TCU, a separate piece of
    hardware in the SoC, provides this functionality.

    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 181bc6b98b109cacd11f9970fef0668532889877
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sun May 30 18:17:57 2021 +0100

    MIPS: cpu-probe: Fix FPU detection on Ingenic JZ4760(B)

    [ Upstream commit fc52f92a653215fbd6bc522ac5311857b335e589 ]

    Ingenic JZ4760 and JZ4760B do have a FPU, but the config registers don't
    report it. Force the FPU detection in case the processor ID match the
    JZ4760(B) one.

    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8853d2ce4e9e96c7b2a9908f752ab2253c99c6ab
Author: Willy Tarreau <w@1wt.eu>
Date:   Sat May 29 13:07:46 2021 +0200

    ipv6: use prandom_u32() for ID generation

    [ Upstream commit 62f20e068ccc50d6ab66fdb72ba90da2b9418c99 ]

    This is a complement to commit aa6dd211e4b1 ("inet: use bigger hash
    table for IP ID generation"), but focusing on some specific aspects
    of IPv6.

    Contary to IPv4, IPv6 only uses packet IDs with fragments, and with a
    minimum MTU of 1280, it's much less easy to force a remote peer to
    produce many fragments to explore its ID sequence. In addition packet
    IDs are 32-bit in IPv6, which further complicates their analysis. On
    the other hand, it is often easier to choose among plenty of possible
    source addresses and partially work around the bigger hash table the
    commit above permits, which leaves IPv6 partially exposed to some
```

```
    possibilities of remote analysis at the risk of weakening some
    protocols like DNS if some IDs can be predicted with a good enough
    probability.

    Given the wide range of permitted IDs, the risk of collision is extremely
    low so there's no need to rely on the positive increment algorithm that
    is shared with the IPv4 code via ip_idents_reserve(). We have a fast
    PRNG, so let's simply call prandom_u32() and be done with it.

    Performance measurements at 10 Gbps couldn't show any difference with
    the previous code, even when using a single core, because due to the
    large fragments, we're limited to only ~930 kpps at 10 Gbps and the cost
    of the random generation is completely offset by other operations and by
    the network transfer time. In addition, this change removes the need to
    update a shared entry in the idents table so it may even end up being
    slightly faster on large scale systems where this matters.

    The risk of at least one collision here is about 1/80 million among
    10 IDs, 1/850k among 100 IDs, and still only 1/8.5k among 1000 IDs,
    which remains very low compared to IPv4 where all IDs are reused
    every 4 to 80ms on a 10 Gbps flow depending on packet sizes.

    Reported-by: Amit Klein <aksecurity@gmail.com>
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Link: https://lore.kernel.org/r/20210529110746.6796-1-w@1wt.eu
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba710baa1cc1b17a0483f7befe03e696efd17292
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Mon May 31 21:58:52 2021 +0800

    virtio-net: Add validation for used length

    [ Upstream commit ad993a95c508417acdeb15244109e009e50d8758 ]

    This adds validation for used length (might come
    from an untrusted device) to avoid data corruption
    or loss.

    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Link: https://lore.kernel.org/r/20210531135852.113-1-xieyongji@bytedance.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8d7221a987416e619af636e32fef946681cf1f02
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Mon May 31 21:56:22 2021 +0800

    drm: bridge: cdns-mhdp8546: Fix PM reference leak in

    [ Upstream commit f674555ee5444c8987dfea0922f1cf6bf0c12847 ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210531135622.3348252-1-
```

```
  yukuai3@huawei.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4dbd10b868b69e300e203513c2f57591f26e6551
  Author: Dmitry Osipenko <digetx@gmail.com>
  Date:   Sun May 16 19:30:35 2021 +0300

      clk: tegra: Ensure that PLLU configuration is applied properly

      [ Upstream commit a7196048cd5168096c2c4f44a3939d7a6dcd06b9 ]

      The PLLU (USB) consists of the PLL configuration itself and configuration
      of the PLLU outputs. The PLLU programming is inconsistent on T30 vs T114,
      where T114 immediately bails out if PLLU is enabled and T30 re-enables
      a potentially already enabled PLL (left after bootloader) and then fully
      reprograms it, which could be unsafe to do. The correct way should be to
      skip enabling of the PLL if it's already enabled and then apply
      configuration to the outputs. This patch doesn't fix any known problems,
      it's a minor improvement.

      Acked-by: Thierry Reding <treding@nvidia.com>
      Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
      Signed-off-by: Thierry Reding <treding@nvidia.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 91789532c24cf7ebc3b2a154ced44e8d2ca85b9b
  Author: Dmitry Osipenko <digetx@gmail.com>
  Date:   Sun May 16 19:30:34 2021 +0300

      clk: tegra: Fix refcounting of gate clocks

      [ Upstream commit c592c8a28f5821e880ac6675781cd8a151b0737c ]

      The refcounting of the gate clocks has a bug causing the enable_refcnt
      to underflow when unused clocks are disabled. This happens because clk
      provider erroneously bumps the refcount if clock is enabled at a boot
      time, which it shouldn't be doing, and it does this only for the gate
      clocks, while peripheral clocks are using the same gate ops and the
      peripheral clocks are missing the initial bump. Hence the refcount of
      the peripheral clocks is 0 when unused clocks are disabled and then the
      counter is decremented further by the gate ops, causing the integer
      underflow.

      Fix this problem by removing the erroneous bump and by implementing the
      disable_unused() callback, which disables the unused gates properly.

      The visible effect of the bug is such that the unused clocks are never
      gated if a loaded kernel module grabs the unused clocks and starts to use
      them. In practice this shouldn't cause any real problems for the drivers
      and boards supported by the kernel today.

      Acked-by: Thierry Reding <treding@nvidia.com>
      Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
      Signed-off-by: Thierry Reding <treding@nvidia.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 9bce65a588d9d3b5fe40db4f415c6afdc0568e0f
  Author: Thierry Reding <treding@nvidia.com>
  Date:   Thu May 27 20:15:46 2021 +0200

      drm/tegra: hub: Fix YUV support

      [ Upstream commit e16efff4e5f490ce34a8c60d9ae7297dca5eb616 ]

      The driver currently exposes several YUV formats but fails to properly
```

```
    program all the registers needed to display such formats. Add the right
    programming sequences so that overlay windows can be used to accelerate
    color format conversions in multimedia playback use-cases.

    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 86e0ca4c6aee48ac7854affa258592b304461d9d
Author: Gioh Kim <gi-oh.kim@cloud.ionos.com>
Date:   Fri May 28 13:30:03 2021 +0200

    RDMA/rtrs: Change MAX_SESS_QUEUE_DEPTH

    [ Upstream commit 3a98ea7041b7d18ac356da64823c2ba2f8391b3e ]

    Max IB immediate data size is 2^28 (MAX_IMM_PAYL_BITS)
    and the minimum chunk size is 4096 (2^12).
    Therefore the maximum sess_queue_depth is 65536 (2^16).

    Link: https://lore.kernel.org/r/20210528113018.52290-6-jinpu.wang@ionos.com
    Signed-off-by: Gioh Kim <gi-oh.kim@ionos.com>
    Signed-off-by: Jack Wang <jinpu.wang@ionos.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f2aad1098936c4abd5a68f7839d57020563d886e
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Thu May 27 18:59:59 2021 +0300

    net: stmmac: the XPCS obscures a potential "PHY not found" error

    [ Upstream commit 4751d2aa321f2828d8c5d2f7ce4ed18a01e47f46 ]

    stmmac_mdio_register() has logic to search for PHYs on the MDIO bus and
    assign them IRQ lines, as well as to set priv->plat->phy_addr.

    If no PHY is found, the "found" variable remains set to 0 and the
    function errors out.

    After the introduction of commit f213bbe8a9d6 ("net: stmmac: Integrate
    it with DesignWare XPCS"), the "found" variable was immediately reused
    for searching for a PCS on the same MDIO bus.

    This can result in 2 types of potential problems (none of them seems to
    be seen on the only Intel system that sets has_xpcs = true, otherwise it
    would have been reported):

    1. If a PCS is found but a PHY is not, then the code happily exits with
       no error. One might say "yes, but this is not possible, because
       of_mdiobus_register will probe a PHY for all MDIO addresses,
       including for the XPCS, so if an XPCS exists, then a PHY certainly
       exists too". Well, that is not true, see intel_mgbe_common_data():

            /* Ensure mdio bus scan skips intel serdes and pcs-xpcs */
            plat->mdio_bus_data->phy_mask = 1 << INTEL_MGBE_ADHOC_ADDR;
            plat->mdio_bus_data->phy_mask |= 1 << INTEL_MGBE_XPCS_ADDR;

    2. A PHY is found but an MDIO device with the XPCS PHY ID isn't, and in
       that case, the error message will be "No PHY found". Confusing.

    Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Link: https://lore.kernel.org/r/20210527155959.3270478-1-olteanv@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 7366ccbd327cc2030b8957df5c4b8a1ac0e1542f
Author: Alex Bee <knaerzche@gmail.com>
Date:   Fri May 28 15:05:52 2021 +0200

    drm: rockchip: add missing registers for RK3066

    [ Upstream commit 742203cd56d150eb7884eb45abb7d9dbc2bdbf04 ]

    Add dither_up, dsp_lut_en and data_blank registers to enable their
    respective functionality for RK3066's VOP.

    While at that also fix .rb_swap and .format registers for all windows,
    which have to be set though RK3066_SYS_CTRL1 register.
    Also remove .scl from win1: Scaling is only supported on the primary
    plane.

    Signed-off-by: Alex Bee <knaerzche@gmail.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210528130554.72191-4-
knaerzche@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20f88209526edf0de3fd9a47bdab0a2c2cb0a57a
Author: Alex Bee <knaerzche@gmail.com>
Date:   Fri May 28 15:05:51 2021 +0200

    drm: rockchip: add missing registers for RK3188

    [ Upstream commit ab64b448a175b8a5a4bd323b8f74758c2574482c ]

    Add dither_up, dsp_lut_en and data_blank registers to enable their
    respective functionality for RK3188's VOP.
    While at that also fix .dsp_blank register which is (only) set with
    BIT24 (same as RK3066)

    Signed-off-by: Alex Bee <knaerzche@gmail.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210528130554.72191-3-
knaerzche@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c77ad9cbf0c70d182d22c26c481958f98dc717b4
Author: Eli Cohen <elic@nvidia.com>
Date:   Mon May 3 14:39:58 2021 +0300

    net/mlx5: Fix lag port remapping logic

    [ Upstream commit 8613641063617c1dfc731b403b3ee4935ef15f87 ]

    Fix the logic so that if both ports netdevices are enabled or disabled,
    use the trivial mapping without swapping.

    If only one of the netdevice's tx is enabled, use it to remap traffic to
    that port.

    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f39de3cc0ff799038aef2987fc70207eccb95289
Author: Huy Nguyen <huyn@nvidia.com>
Date:   Tue Dec 15 10:58:54 2020 -0600

    net/mlx5e: IPsec/rep_tc: Fix rep_tc_update_skb drops IPsec packet
```

```
    [ Upstream commit c07274ab1ab2c38fb128e32643c22c89cb319384 ]

    rep_tc copy REG_C1 to REG_B. IPsec crypto utilizes the whole REG_B
    register with BIT31 as IPsec marker. rep_tc_update_skb drops
    IPsec because it thought REG_B contains bad value.

    In previous patch, BIT 31 of REG_C1 is reserved for IPsec.
    Skip the rep_tc_update_skb if BIT31 of REG_B is set.

    Signed-off-by: Huy Nguyen <huyn@nvidia.com>
    Signed-off-by: Raed Salem <raeds@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba2d94789a1f2ee146276dc3dd00f36d348387ad
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Wed May 26 23:47:02 2021 -0400

    drm/amdgpu/swsmu/aldebaran: fix check in is_dpm_running

    [ Upstream commit dd1d82c04e111b5a864638ede8965db2fe6d8653 ]

    If smu_cmn_get_enabled_mask() fails, return false to be
    consistent with other asics.

    Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: Lee Jones <lee.jones@linaro.org>
    Reviewed-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 93ed1b0210ea130d61d13b0b95297a7d53e08110
Author: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
Date:   Mon May 10 10:19:18 2021 -0400

    drm/amd/display: fix odm scaling

    [ Upstream commit 6566cae7aef30da8833f1fa0eb854baf33b96676 ]

    There are two issues with scaling calculations, odm recout
    calculation and matching viewport to actual recout.

    This change fixes both issues. Odm recout calculation via
    special casing and viewport matching issue by reworking
    the viewport calcualtion to use scaling ratios and recout
    to derrive the required offset and size.

    Signed-off-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Reviewed-by: Jun Lei <Jun.Lei@amd.com>
    Acked-by: Qingqing Zhuo <qingqing.zhuo@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 669f0f6f83d53fedfe1ad35c571e183a3a8a8416
Author: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
Date:   Thu May 27 13:36:38 2021 +0900

    clk: renesas: r8a77995: Add ZA2 clock

    [ Upstream commit 790c06cc5df263cdaff748670cc65958c81b0951 ]

    R-Car D3 ZA2 clock is from PLL0D3 or S0,
    and it can be controlled by ZA2CKCR.
```

```
    It is needed for R-Car Sound, but is not used so far.
    Using default settings is very enough at this point.
    This patch adds it by DEF_FIXED().

    Signed-off-by: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
    Link: https://lore.kernel.org/r/87pmxclrmy.wl-kuninori.morimoto.gx@renesas.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3831144772502d86d8c22b3f1390a19bb462c0e0
Author: Zou Wei <zou_wei@huawei.com>
Date:   Mon May 24 15:21:02 2021 +0800

    drm/bridge: cdns: Fix PM reference leak in cdns_dsi_transfer()

    [ Upstream commit 33f90f27e1c5ccd648d3e78a1c28be9ee8791cf1 ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/1621840862-106024-1-git-send-email-
zou_wei@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4432f3694ee53a8091467d7949ca96eb9a51267a
Author: Jesse Brandeburg <jesse.brandeburg@intel.com>
Date:   Thu Mar 25 17:38:29 2021 -0700

    igb: fix assignment on big endian machines

    [ Upstream commit b514958dd1a3bd57638b0e63b8e5152b1960e6aa ]

    The igb driver was trying hard to be sparse correct, but somehow
    ended up converting a variable into little endian order and then
    tries to OR something with it.

    A much plainer way of doing things is to leave all variables and
    OR operations in CPU (non-endian) mode, and then convert to
    little endian only once, which is what this change does.

    This probably fixes a bug that might have been seen only on
    big endian systems.

    Signed-off-by: Jesse Brandeburg <jesse.brandeburg@intel.com>
    Tested-by: Dave Switzer <david.switzer@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 003d37fa8c70d8067d73ccad86241e839919fd9b
Author: Jesse Brandeburg <jesse.brandeburg@intel.com>
Date:   Thu Mar 25 17:38:28 2021 -0700

    igb: handle vlan types with checker enabled

    [ Upstream commit c7cbfb028b95360403d579c47aaaeef1ff140964 ]

    The sparse build (C=2) finds some issues with how the driver
    dealt with the (very difficult) hardware that in some generations
    uses little-endian, and in others uses big endian, for the VLAN
```

```
    field. The code as written picks __le16 as a type and for some
    hardware revisions we override it to __be16 as done in this
    patch. This impacted the VF driver as well so fix it there too.

    Also change the vlan_tci assignment to override the sparse
    warning without changing functionality.

    Signed-off-by: Jesse Brandeburg <jesse.brandeburg@intel.com>
    Tested-by: Dave Switzer <david.switzer@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd658e324c34a5c760a75b52f02181318003c18b
Author: Jesse Brandeburg <jesse.brandeburg@intel.com>
Date:   Thu Mar 25 17:38:24 2021 -0700

    e100: handle eeprom as little endian

    [ Upstream commit d4ef55288aa2e1b76033717242728ac98ddc4721 ]

    Sparse tool was warning on some implicit conversions from
    little endian data read from the EEPROM on the e100 cards.

    Fix these by being explicit about the conversions using
    le16_to_cpu().

    Signed-off-by: Jesse Brandeburg <jesse.brandeburg@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 38d70d239c7c9a51f0bde26608ac77ff0356d6df
Author: Zou Wei <zou_wei@huawei.com>
Date:   Mon May 24 15:20:54 2021 +0800

    drm/vc4: hdmi: Fix PM reference leak in vc4_hdmi_encoder_pre_crtc_co()

    [ Upstream commit 5e4322a8b266bc9f5ee7ea4895f661c01dbd7cb3 ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/1621840854-105978-1-git-send-email-
zou_wei@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 62f616078e9f6338a8fb08c4f8f3c48609a12da8
Author: Mateusz Kwiatkowski <kfyatek+publicgit@gmail.com>
Date:   Thu May 20 17:03:41 2021 +0200

    drm/vc4: Fix clock source for VEC PixelValve on BCM2711

    [ Upstream commit fc7a8abcee2225d6279ff785d33e24d70c738c6e ]

    On the BCM2711 (Raspberry Pi 4), the VEC is actually connected to
    output 2 of pixelvalve3.

    NOTE: This contradicts the Broadcom docs, but has been empirically
    tested and confirmed by Raspberry Pi firmware devs.

    Signed-off-by: Mateusz Kwiatkowski <kfyatek+publicgit@gmail.com>
```

```
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210520150344.273900-2-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d0f4887e7429ac4a56dfa2844d72b67fe9a70c8
Author: Feifei Xu <Feifei.Xu@amd.com>
Date:   Thu May 20 17:46:03 2021 +0800

    drm/amd/pm: fix return value in aldebaran_set_mp1_state()

    [ Upstream commit 5051cb794ac5d92154e186d87cdc12cba613f4f6 ]

    For default cases,we should return 0. Otherwise resume will
    abort because of the wrong return value.

    Signed-off-by: Feifei Xu <Feifei.Xu@amd.com>
    Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8722275b41d5127048e1422a8a1b6370b4878533
Author: YueHaibing <yuehaibing@huawei.com>
Date:   Wed May 19 10:47:04 2021 +0800

    net: xilinx_emaclite: Do not print real IOMEM pointer

    [ Upstream commit d0d62baa7f505bd4c59cd169692ff07ec49dde37 ]

    Printing kernel pointers is discouraged because they might leak kernel
    memory layout.  This fixes smatch warning:

    drivers/net/ethernet/xilinx/xilinx_emaclite.c:1191 xemaclite_of_probe() warn:
     argument 4 to %08lX specifier is cast from pointer

    Signed-off-by: YueHaibing <yuehaibing@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5150877e4d99f85057a458daac7cd7c01005d5c6
Author: Arturo Giusti <koredump@protonmail.com>
Date:   Tue May 18 12:34:57 2021 +0200

    udf: Fix NULL pointer dereference in udf_symlink function

    [ Upstream commit fa236c2b2d4436d9f19ee4e5d5924e90ffd7bb43 ]

    In function udf_symlink, epos.bh is assigned with the value returned
    by udf_tgetblk. The function udf_tgetblk is defined in udf/misc.c
    and returns the value of sb_getblk function that could be NULL.
    Then, epos.bh is used without any check, causing a possible
    NULL pointer dereference when sb_getblk fails.

    This fix adds a check to validate the value of epos.bh.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=213083
    Signed-off-by: Arturo Giusti <koredump@protonmail.com>
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50d7e03ad487cc45fc85164a299b945a41756ac0
Author: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
Date:   Wed May 19 10:14:07 2021 -0400
```

```
    drm/sched: Avoid data corruptions

    [ Upstream commit 0b10ab80695d61422337ede6ff496552d8ace99d ]

    Wait for all dependencies of a job  to complete before
    killing it to avoid data corruptions.

    Signed-off-by: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210519141407.88444-1-
andrey.grodzovsky@amd.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit baf802cd27cd95fb027511b8bf5d1fd6fdf80907
Author: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
Date:   Wed May 12 10:26:45 2021 -0400

    drm/scheduler: Fix hang when sched_entity released

    [ Upstream commit c61cdbdbffc169dc7f1e6fe94dfffaf574fe672a ]

    Problem: If scheduler is already stopped by the time sched_entity
    is released and entity's job_queue not empty I encountred
    a hang in drm_sched_entity_flush. This is because drm_sched_entity_is_idle
    never becomes false.

    Fix: In drm_sched_fini detach all sched_entities from the
    scheduler's run queues. This will satisfy drm_sched_entity_is_idle.
    Also wakeup all those processes stuck in sched_entity flushing
    as the scheduler main thread which wakes them up is stopped by now.

    v2:
    Reverse order of drm_sched_rq_remove_entity and marking
    s_entity as stopped to prevent reinserion back to rq due
    to race.

    v3:
    Drop drm_sched_rq_remove_entity, only modify entity->stopped
    and check for it in drm_sched_entity_is_idle

    Signed-off-by: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210512142648.666476-14-
andrey.grodzovsky@amd.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cef0b8eab2310b51589b41ea839245dce6c95ca7
Author: Bixuan Cui <cuibixuan@huawei.com>
Date:   Sat May 8 11:15:02 2021 +0800

    pinctrl: equilibrium: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit d7f444499d6faf9a6ae3b27ec094109528d2b9a7 ]

    This patch adds missing MODULE_DEVICE_TABLE definition which generates
    correct modalias for automatic loading of this driver when it is built
    as an external module.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Bixuan Cui <cuibixuan@huawei.com>
    Link: https://lore.kernel.org/r/20210508031502.53637-1-cuibixuan@huawei.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 248fc2633293a9eb9c1686222170b52e76c80baa
```

```
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Wed May 19 14:16:27 2021 +0000

    net: ethernet: ixp4xx: Fix return value check in ixp4xx_eth_probe()

    [ Upstream commit 20e76d3d044d936998617f8acd7e77bebd9ca703 ]

    In case of error, the function mdiobus_get_phy() returns NULL
    pointer not ERR_PTR(). The IS_ERR() test in the return value
    check should be replaced with NULL test.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b777772879c88fe27984d3677c90a25feb3a2f2
Author: Davide Caratti <dcaratti@redhat.com>
Date:   Wed May 19 15:17:21 2021 +0200

    net/sched: cls_api: increase max_reclassify_loop

    [ Upstream commit 05ff8435e50569a0a6b95e5ceaea43696e8827ab ]

    modern userspace applications, like OVN, can configure the TC datapath to
    "recirculate" packets several times. If more than 4 "recirculation" rules
    are configured, packets can be dropped by __tcf_classify().
    Changing the maximum number of reclassifications (from 4 to 16) should be
    sufficient to prevent drops in most use cases, and guard against loops at
    the same time.

    Signed-off-by: Davide Caratti <dcaratti@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f20ce74aff8de015608cb867aa06367315ae8d4
Author: Vladimir Oltean <olteanv@gmail.com>
Date:   Tue May 18 20:49:24 2021 +0300

    net: mdio: provide shim implementation of devm_of_mdiobus_register

    [ Upstream commit 86544c3de6a2185409c5a3d02f674ea223a14217 ]

    Similar to the way in which of_mdiobus_register() has a fallback to the
    non-DT based mdiobus_register() when CONFIG_OF is not set, we can create
    a shim for the device-managed devm_of_mdiobus_register() which calls
    devm_mdiobus_register() and discards the struct device_node *.

    In particular, this solves a build issue with the qca8k DSA driver which
    uses devm_of_mdiobus_register and can be compiled without CONFIG_OF.

    Reported-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Vladimir Oltean <olteanv@gmail.com>
    Acked-by: Randy Dunlap <rdunlap@infradead.org> # build-tested
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Reviewed-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5cb3b25c7f5318199767f071b780ca56a8a4141e
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Mon May 17 16:49:12 2021 +0800

    drm/virtio: Fix double free on probe failure
```

```
    [ Upstream commit cec7f1774605a5ef47c134af62afe7c75c30b0ee ]

    The virtio_gpu_init() will free vgdev and vgdev->vbufs on failure.
    But such failure will be caught by virtio_gpu_probe() and then
    virtio_gpu_release() will be called to do some cleanup which
    will free vgdev and vgdev->vbufs again. So let's set dev->dev_private
    to NULL to avoid double free.

    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Link: http://patchwork.freedesktop.org/patch/msgid/20210517084913.403-2-
xieyongji@bytedance.com
    Signed-off-by: Gerd Hoffmann <kraxel@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5ac1c5fcffd0cf61718cfa4459847b948f9c6978
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Mon May 17 15:15:45 2021 +0300

    reiserfs: add check for invalid 1st journal block

    [ Upstream commit a149127be52fa7eaf5b3681a0317a2bbb772d5a9 ]

    syzbot reported divide error in reiserfs.
    The problem was in incorrect journal 1st block.

    Syzbot's reproducer manualy generated wrong superblock
    with incorrect 1st block. In journal_init() wasn't
    any checks about this particular case.

    For example, if 1st journal block is before superblock
    1st block, it can cause zeroing important superblock members
    in do_journal_end().

    Link: https://lore.kernel.org/r/20210517121545.29645-1-paskripkin@gmail.com
    Reported-by: syzbot+0ba9909df31c6a36974d@syzkaller.appspotmail.com
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Jan Kara <jack@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8f63c943700e2f77909fe628b83583f99d9e9d2
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 14:45:55 2021 +0800

    drm/bridge: lt9611: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 8d0b1fe81e18eb66a2d4406386760795fe0d77d9 ]

    This patch adds missing MODULE_DEVICE_TABLE definition which generates
    correct modalias for automatic loading of this driver when it is built
    as an external module.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/1620801955-19188-1-git-send-email-
zou_wei@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3828625dd18d1fb2fba62015791ad0c6ddb9b402
Author: Ansuel Smith <ansuelsmth@gmail.com>
Date:   Fri May 14 23:03:50 2021 +0200

    net: mdio: ipq8064: add regmap config to disable REGCACHE
```

```
    [ Upstream commit b097bea10215315e8ee17f88b4c1bbb521b1878c ]

    mdio drivers should not use REGCHACHE. Also disable locking since it's
    handled by the mdio users and regmap is always accessed atomically.

    Signed-off-by: Ansuel Smith <ansuelsmth@gmail.com>
    Reviewed-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8f5cc4dc1e2a2f9e77bd8bee66e7940c2222adae
Author: Nicolas Boichat <drinkcat@chromium.org>
Date:   Wed Apr 21 13:28:54 2021 +0800

    drm/panfrost: devfreq: Disable devfreq when num_supplies > 1

    [ Upstream commit 09da3191827f2fd326205fb58881838e6ea36fb7 ]

    GPUs with more than a single regulator (e.g. G72 on MT8183) will
    require platform-specific handling for devfreq, for 2 reasons:
     1. The opp core (drivers/opp/core.c:_generic_set_opp_regulator)
        does not support multiple regulators, so we'll need custom
        handlers.
     2. Generally, platforms with 2 regulators have platform-specific
        constraints on how the voltages should be set (e.g.
        minimum/maximum voltage difference between them), so we
        should not just create generic handlers that simply
        change the voltages without taking care of those constraints.

    Disable devfreq for now on those GPUs.

    Signed-off-by: Nicolas Boichat <drinkcat@chromium.org>
    Reviewed-by: Tomeu Vizoso <tomeu.vizoso@collabora.com>
    Reviewed-by: Steven Price <steven.price@arm.com>
    Signed-off-by: Steven Price <steven.price@arm.com>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210421132841.v13.3.I3af068abe30c9c85cabc4486
385c52e56527a509@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e049b85f7dcc666921e7513bf22b70d9462e624d
Author: Wang Li <wangli74@huawei.com>
Date:   Sat Apr 10 03:48:41 2021 +0000

    drm/mediatek: Fix PM reference leak in mtk_crtc_ddp_hw_init()

    [ Upstream commit 69777e6ca396f0a7e1baff40fcad4a9d3d445b7a ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Li <wangli74@huawei.com>
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 609ee98affb06f1013331a57efc281372609fa0f
Author: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
Date:   Wed May 12 23:43:24 2021 +0200

    net: Treat __napi_schedule_irqoff() as __napi_schedule() on PREEMPT_RT

    [ Upstream commit 8380c81d5c4fced6f4397795a5ae65758272bbfd ]
```

```
    __napi_schedule_irqoff() is an optimized version of __napi_schedule()
    which can be used where it is known that interrupts are disabled,
    e.g. in interrupt-handlers, spin_lock_irq() sections or hrtimer
    callbacks.

    On PREEMPT_RT enabled kernels this assumptions is not true. Force-
    threaded interrupt handlers and spinlocks are not disabling interrupts
    and the NAPI hrtimer callback is forced into softirq context which runs
    with interrupts enabled as well.

    Chasing all usage sites of __napi_schedule_irqoff() is a whack-a-mole
    game so make __napi_schedule_irqoff() invoke __napi_schedule() for
    PREEMPT_RT kernels.

    The callers of ____napi_schedule() in the networking core have been
    audited and are correct on PREEMPT_RT kernels as well.

    Reported-by: Juri Lelli <juri.lelli@redhat.com>
    Signed-off-by: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
    Reviewed-by: Thomas Gleixner <tglx@linutronix.de>
    Reviewed-by: Juri Lelli <juri.lelli@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b991df8881088448cb223e769e37cab8dd40706
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 15:00:24 2021 +0800

    atm: nicstar: Fix possible use-after-free in nicstar_cleanup()

    [ Upstream commit 34e7434ba4e97f4b85c1423a59b2922ba7dff2ea ]

    This module's remove path calls del_timer(). However, that function
    does not wait until the timer handler finishes. This means that the
    timer handler may still be running after the driver's remove function
    has finished, which would result in a use-after-free.

    Fix by calling del_timer_sync(), which makes sure the timer handler
    has finished, and unable to re-schedule itself.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed7c3739d0a07e2ec3ccbffe7e93cea01c438cda
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 14:58:53 2021 +0800

    mISDN: fix possible use-after-free in HFC_cleanup()

    [ Upstream commit 009fc857c5f6fda81f2f7dd851b2d54193a8e733 ]

    This module's remove path calls del_timer(). However, that function
    does not wait until the timer handler finishes. This means that the
    timer handler may still be running after the driver's remove function
    has finished, which would result in a use-after-free.

    Fix by calling del_timer_sync(), which makes sure the timer handler
    has finished, and unable to re-schedule itself.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
```

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

commit d1fb12412874c94ad037e11d0ecdd1140a439297
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 14:53:36 2021 +0800

    atm: iphase: fix possible use-after-free in ia_module_exit()

    [ Upstream commit 1c72e6ab66b9598cac741ed397438a52065a8f1f ]

    This module's remove path calls del_timer(). However, that function
    does not wait until the timer handler finishes. This means that the
    timer handler may still be running after the driver's remove function
    has finished, which would result in a use-after-free.

    Fix by calling del_timer_sync(), which makes sure the timer handler
    has finished, and unable to re-schedule itself.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 72bdf3c6b059b2bab3974dea9882cf381f12d7a0
Author: Bibo Mao <maobibo@loongson.cn>
Date:   Mon Jun 29 21:15:32 2020 +0800

    hugetlb: clear huge pte during flush function on mips platform

    [ Upstream commit 33ae8f801ad8bec48e886d368739feb2816478f2 ]

    If multiple threads are accessing the same huge page at the same
    time, hugetlb_cow will be called if one thread write the COW huge
    page. And function huge_ptep_clear_flush is called to notify other
    threads to clear the huge pte tlb entry. The other threads clear
    the huge pte tlb entry and reload it from page table, the reload
    huge pte entry may be old.

    This patch fixes this issue on mips platform, and it clears huge
    pte entry before notifying other threads to flush current huge
    page entry, it is similar with other architectures.

    Signed-off-by: Bibo Mao <maobibo@loongson.cn>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d2958a2e6d9246a41e649bafd786018add603c2f
Author: Youling Tang <tangyouling@loongson.cn>
Date:   Thu May 6 10:02:50 2021 +0800

    MIPS: Loongson64: Fix build error 'secondary_kexec_args' undeclared under !SMP

    [ Upstream commit 6a73022ee3fdf7e60f2ba0a3a835dd421c05b5b5 ]

    On the Loongson64 platform, if CONFIG_SMP is not set, the following build
    error will occur:
    arch/mips/loongson64/reset.c:133:2: error:'secondary_kexec_args' undeclared

    Because the definition and declaration of secondary_kexec_args are in the
    CONFIG_SMP, the secondary_kexec_args variable should be used in CONFIG_SMP.

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Youling Tang <tangyouling@loongson.cn>
    Acked-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
```

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

commit 6102f34e462285174b590c50a37443070dfda42e
Author: Dinghao Liu <dinghao.liu@zju.edu.cn>
Date:   Thu Apr 15 15:33:38 2021 +0800

    clk: renesas: rcar-usb2-clock-sel: Fix error handling in .probe()

    [ Upstream commit a20a40a8bbc2cf4b29d7248ea31e974e9103dd7f ]

    The error handling paths after pm_runtime_get_sync() have no refcount
    decrement, which leads to refcount leak.

    Signed-off-by: Dinghao Liu <dinghao.liu@zju.edu.cn>
    Link: https://lore.kernel.org/r/20210415073338.22287-1-dinghao.liu@zju.edu.cn
    [geert: Remove now unused variable priv]
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5ec05aa8160bc59df13e4aa4c758b49aa035be0e
Author: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
Date:   Mon Apr 19 17:50:53 2021 -0400

    drm/amd/display: fix use_max_lb flag for 420 pixel formats

    [ Upstream commit 8809a7a4afe90ad9ffb42f72154d27e7c47551ae ]

    Right now the flag simply selects memory config 0 when flag is true
    however 420 modes benefit more from memory config 3.

    Signed-off-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Reviewed-by: Aric Cyr <Aric.Cyr@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57092f4f803b31cb1fddfec2bb33196967c6e946
Author: Ilya Bakoulin <Ilya.Bakoulin@amd.com>
Date:   Mon Apr 26 14:27:38 2021 -0400

    drm/amd/display: Fix clock table filling logic

    [ Upstream commit c31bef1cb1203b26f901a511a3246204cfaf8a57 ]

    [Why]
    Currently, the code that fills the clock table can miss filling
    information about some of the higher voltage states advertised
    by the SMU. This, in turn, may cause some of the higher pixel clock
    modes (e.g. 8k60) to fail validation.

    [How]
    Fill the table with one entry per DCFCLK level instead of one entry
    per FCLK level. This is needed because the maximum FCLK does not
    necessarily need maximum voltage, whereas DCFCLK values from SMU
    cover the full voltage range.

    Signed-off-by: Ilya Bakoulin <Ilya.Bakoulin@amd.com>
    Reviewed-by: Dmytro Laktyushkin <Dmytro.Laktyushkin@amd.com>
    Acked-by: Stylon Wang <stylon.wang@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3eb50ab2c5ac3e25c4ea056491474baac78a4fb3
```

```
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Tue May 4 11:00:42 2021 -0400

    drm/amdgpu: change the default timeout for kernel compute queues

    [ Upstream commit 67387dfe0f6630f2d4f412ce77debec23a49db7a ]

    Change to 60s.  This matches what we already do in virtualization.
    Infinite timeout can lead to deadlocks in the kernel.

    Reviewed-by: Christian König <christian.koenig@amd.com>
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9c00298cf07e6628f8fb7b6a33e473bb05ae45c
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Mon May 10 19:39:30 2021 +0300

    net: pch_gbe: Use proper accessors to BE data in pch_ptp_match()

    [ Upstream commit 443ef39b499cc9c6635f83238101f1bb923e9326 ]

    Sparse is not happy about handling of strict types in pch_ptp_match():

      .../pch_gbe_main.c:158:33: warning: incorrect type in argument 2 (different base
types)
      .../pch_gbe_main.c:158:33:    expected unsigned short [usertype] uid_hi
      .../pch_gbe_main.c:158:33:    got restricted __be16 [usertype]
      .../pch_gbe_main.c:158:45: warning: incorrect type in argument 3 (different base
types)
      .../pch_gbe_main.c:158:45:    expected unsigned int [usertype] uid_lo
      .../pch_gbe_main.c:158:45:    got restricted __be32 [usertype]
      .../pch_gbe_main.c:158:56: warning: incorrect type in argument 4 (different base
types)
      .../pch_gbe_main.c:158:56:    expected unsigned short [usertype] seqid
      .../pch_gbe_main.c:158:56:    got restricted __be16 [usertype]

    Fix that by switching to use proper accessors to BE data.

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Flavio Suligoi <f.suligoi@asem.it>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae1864f13cce523b1e7811f3ec4ff931abcd965c
Author: Sebastian Reichel <sebastian.reichel@collabora.com>
Date:   Thu Apr 29 00:29:50 2021 +0200

    drm/imx: Add 8 pixel alignment fix

    [ Upstream commit 94dfec48fca756cef90263a03e81f24dae24a5c6 ]

    Some standard resolutions like 1366x768 do not work properly with
    i.MX6 SoCs, since the horizontal resolution needs to be aligned
    to 8 pixels (so 1360x768 or 1368x768 would work).

    This patch allocates framebuffers allocated to 8 pixels. The extra
    time required to send the extra pixels are removed from the blank
    time. In order to expose the correct display size to userspace,
    the stride is increased without increasing the width.

    Without this patch systems with this display resolution hang
    indefinitely during boot up.
```

```
    Suggested-by: Boris Brezillon <boris.brezillon@collabora.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Link: https://lore.kernel.org/r/20210428222953.235280-3-sebastian.reichel@collabora.com
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 925a6b206bde1df6f7bba951b0d006c921c36c8f
Author: Liu Ying <victor.liu@nxp.com>
Date:   Fri Apr 23 17:26:41 2021 +0800

    drm/bridge: nwl-dsi: Force a full modeset when crtc_state->active is changed to be true

    [ Upstream commit 3afb2a28fa2404d11cce1956a003f2aaca4da421 ]

    This patch replaces ->mode_fixup() with ->atomic_check() so that
    a full modeset can be requested from there when crtc_state->active
    is changed to be true(which implies only connector's DPMS is brought
    out of "Off" status, though not necessarily).  Bridge functions are
    added or changed to accommodate the ->atomic_check() callback.  That
    full modeset is needed by the up-coming patch which gets MIPI DSI
    controller and PHY ready in ->mode_set(), because it makes sure
    ->mode_set() and ->atomic_disable() are called in pairs.

    Cc: Andrzej Hajda <a.hajda@samsung.com>
    Cc: Neil Armstrong <narmstrong@baylibre.com>
    Cc: Robert Foss <robert.foss@linaro.org>
    Cc: Laurent Pinchart <Laurent.pinchart@ideasonboard.com>
    Cc: Jonas Karlman <jonas@kwiboo.se>
    Cc: Jernej Skrabec <jernej.skrabec@siol.net>
    Cc: David Airlie <airlied@linux.ie>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Cc: Guido Günther <agx@sigxcpu.org>
    Cc: Robert Chiras <robert.chiras@nxp.com>
    Cc: NXP Linux Team <linux-imx@nxp.com>
    Signed-off-by: Liu Ying <victor.liu@nxp.com>
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/1619170003-4817-2-git-send-email-
victor.liu@nxp.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0cf0cb019cd5ea9a09b23816b8d2b1707b813f88
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Wed Apr 21 13:18:03 2021 +0300

    drm/vc4: fix argument ordering in vc4_crtc_get_margins()

    [ Upstream commit e590c2b03a6143ba93ddad306bc9eaafa838c020 ]

    Cppcheck complains that the declaration doesn't match the function
    definition.  Obviously "left" should come before "right".  The caller
    and the function implementation are done this way, it's just the
    declaration which is wrong so this doesn't affect runtime.

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/YH/720FD978TPhHp@mwanda
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e684e0ee1153f40509926ead0da8007e8061d9ce
Author: Jack Zhang <Jack.Zhang1@amd.com>
Date:   Tue Apr 27 17:08:47 2021 +0800
```

```
    drm/amd/amdgpu/sriov disable all ip hw status by default

    [ Upstream commit 95ea3dbc4e9548d35ab6fbf67675cef8c293e2f5 ]

    Disable all ip's hw status to false before any hw_init.
    Only set it to true until its hw_init is executed.

    The old 5.9 branch has this change but somehow the 5.11 kernrel does
    not have this fix.

    Without this change, sriov tdr have gfx IB test fail.

    Signed-off-by: Jack Zhang <Jack.Zhang1@amd.com>
    Review-by: Emily Deng <Emily.Deng@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eeb44b638475bcccd9b8da6a4a2274181431453e
Author: Chris Park <Chris.Park@amd.com>
Date:   Tue Apr 13 22:25:23 2021 -0400

    drm/amd/display: Fix BSOD with NULL check

    [ Upstream commit b2d4b9f72fb14c1e6e4f0128964a84539a72d831 ]

    [Why]
    CLK mgr is null for server settings.

    [How]
    Guard the function with NULL check.

    Signed-off-by: Chris Park <Chris.Park@amd.com>
    Reviewed-by: Nicholas Kazlauskas <Nicholas.Kazlauskas@amd.com>
    Acked-by: Wayne Lin <waynelin@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 3728601cd68e3a43ecb49a394cd0056b0ecfba7c
Author: Lewis Huang <Lewis.Huang@amd.com>
Date:   Tue Apr 13 07:02:28 2021 +0800

    drm/amd/display: Revert wait vblank on update dpp clock

    [ Upstream commit d5433a9f692f57c814286f8af2746c567ef79fc8 ]

    [Why]
    This change only fix dpp clock switch to lower case.
    New solution later can fix both case, which is "dc: skip
    program clock when allow seamless boot"

    [How]
    This reverts commit "dc: wait vblank when stream enabled
    and update dpp clock"

    Signed-off-by: Lewis Huang <Lewis.Huang@amd.com>
    Reviewed-by: Eric Yang <eric.yang2@amd.com>
    Acked-by: Wayne Lin <waynelin@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 73893fe5217804bd151ecd346fece25807bb20da
Author: Brandon Syu <Brandon.Syu@amd.com>
Date:   Fri Apr 9 14:47:46 2021 +0800
```

```
    drm/amd/display: fix HDCP reset sequence on reinitialize

    [ Upstream commit 99c248c41c2199bd34232ce8e729d18c4b343b64 ]

    [why]
    When setup is called after hdcp has already setup,
    it would cause to disable HDCP flow won’t execute.

    [how]
    Don't clean up hdcp content to be 0.

    Signed-off-by: Brandon Syu <Brandon.Syu@amd.com>
    Reviewed-by: Wenjing Liu <Wenjing.Liu@amd.com>
    Acked-by: Wayne Lin <waynelin@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 fc2e7aeedba2b7293ba04a24e6edaf57cbd1b422
Author: KuoHsiang Chou <kuohsiang_chou@aspeedtech.com>
Date:   Wed Apr 21 16:58:59 2021 +0800

    drm/ast: Fixed CVE for DP501

    [ Upstream commit ba4e0339a6a33e2ba341703ce14ae8ca203cb2f1 ]

    [Bug][DP501]
    If ASPEED P2A (PCI to AHB) bridge is disabled and disallowed for
    CVE_2019_6260 item3, and then the monitor's EDID is unable read through
    Parade DP501.
    The reason is the DP501's FW is mapped to BMC addressing space rather
    than Host addressing space.
    The resolution is that using "pci_iomap_range()" maps to DP501's FW that
    stored on the end of FB (Frame Buffer).
    In this case, FrameBuffer reserves the last 2MB used for the image of
    DP501.

    Signed-off-by: KuoHsiang Chou <kuohsiang_chou@aspeedtech.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210421085859.17761-1-
kuohsiang_chou@aspeedtech.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d5d89dd11f6fe8e97e3028fddc3b4a14f8d93892
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Thu Apr 15 13:00:39 2021 +0200

    drm/zte: Don't select DRM_KMS_FB_HELPER

    [ Upstream commit a50e74bec1d17e95275909660c6b43ffe11ebcf0 ]

    Selecting DRM_FBDEV_EMULATION will include the correct settings for
    fbdev emulation. Drivers should not override this.

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210415110040.23525-4-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 992b191b179c49250e64f6ec4472abad50a21742
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Thu Apr 15 13:00:38 2021 +0200
```

```
    drm/mxsfb: Don't select DRM_KMS_FB_HELPER

    [ Upstream commit 13b29cc3a722c2c0bc9ab9f72f9047d55d08a2f9 ]

    Selecting DRM_FBDEV_EMULATION will include the correct settings for
    fbdev emulation. Drivers should not override this.

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Acked-by: Stefan Agner <stefan@agner.ch>
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210415110040.23525-3-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```