```
commit 2104f927ad5e7450da588fef81f06dccbfed484e
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Thu Jan 27 12:03:05 2022 +0100

    Linux 5.16.3

    Link: https://lore.kernel.org/r/20220124184125.121143506@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Zan Aziz <zanaziz313@gmail.com>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Rudi Heitbaum <rudi@heitbaum.com>
    Tested-by: Ron Economos <re@w6rz.net>
    Link: https://lore.kernel.org/r/20220125155447.179130255@linuxfoundation.org
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Rudi Heitbaum <rudi@heitbaum.com>
    Tested-by: Salvatore Bonaccorso <carnil@debian.org>
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: <re@w6rz.net>
    Tested-by: Jeffrin Jose T  <jeffrin@rajagiritech.edu.in>
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Tested-by: Scott Bruce <smbruce@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7eef33922335bc7fb97f63c786e19ef2086dafe1
Author: Mauro Carvalho Chehab <mchehab@kernel.org>
Date:   Thu Jan 6 01:41:02 2022 +0100

    scripts: sphinx-pre-install: Fix ctex support on Debian

    commit 87d6576ddf8ac25f36597bc93ca17f6628289c16 upstream.

    The name of the package with ctexhook.sty is different on
    Debian/Ubuntu.

    Reported-by: Akira Yokosawa <akiyks@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Tested-by: Akira Yokosawa <akiyks@gmail.com>
    Link:
https://lore.kernel.org/r/63882425609a2820fac78f5e94620abeb7ed5f6f.1641429634.git.mchehab@ke
rnel.org
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7003566eaf5eae9d28c46f798b7772a8eb4c0ad6
Author: Mauro Carvalho Chehab <mchehab@kernel.org>
Date:   Mon Jan 3 22:01:57 2022 +0100

    scripts: sphinx-pre-install: add required ctex dependency

    commit 7baab965896eaeea60a54b8fe742feea2f79060f upstream.

    After a change meant to fix support for oriental characters
    (Chinese, Japanese, Korean), ctex stylesheet is now a requirement
    for PDF output.

    Reported-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Link:
https://lore.kernel.org/r/165aa6167f21e3892a6e308688c93c756e94f4e0.1641243581.git.mchehab@ke
rnel.org
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
```

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

commit dad470d4a936b7c9e83b630329b7a85e0aa959ff
Author: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
Date:   Tue Nov 23 19:16:06 2021 +0200

    ASoC: SOF: handle paused streams during system suspend

    commit 96da174024b9c63bd5d3358668d0bc12677be877 upstream.

    During system suspend, paused streams do not get suspended.
    Therefore, we need to explicitly free these PCMs in the DSP
    and free the associated DAPM widgets so that they can be set
    up again during resume.

    Fixes: 5fcdbb2d45df ("ASoC: SOF: Add support for dynamic pipelines")
    Signed-off-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Reviewed-by: Paul Olaru <paul.olaru@oss.nxp.com>
    Reviewed-by: Bard Liao <bard.liao@intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20211123171606.129350-3-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b4ef481a7f6ef5770eeca2930bd4db664a8c5287
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Tue Nov 23 19:16:04 2021 +0200

    ASoC: SOF: sof-audio: setup sched widgets during pipeline complete step

    commit 01429183f479c54c1b5d15453a8ce574ea43e525 upstream.

    Older firmware prior to ABI 3.19 has a dependency where the scheduler
    widgets need to be setup last. Moving the call to sof_widget_setup()
    before the pipeline_complete() call also helps remove the need for the
    'reverse' direction when walking through the widget list - this was
    only working because of the topology macros but the topology does not
    require any order.

    Fixes: 5fcdbb2d45df ("ASoC: SOF: Add support for dynamic pipelines")
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20211123171606.129350-1-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f2bb2a32c8eddab4d4cce31f72e419f31add7b3d
Author: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
Date:   Fri Nov 19 21:26:18 2021 +0200

    ASoC: SOF: free widgets in sof_tear_down_pipelines() for static pipelines

    commit b2ebcf42a48f4560862bb811f3268767d17ebdcd upstream.

    Free widgets for static pipelines in sof_tear_down_pipelines().
    But this feature is unavailable in older firmware with ABI < 3.19.
    Just reset widget use_count's for this case. This would ensure that
    the secondary cores enabled required for topology setup are powered
    down properly before the primary core is powered off during
    system suspend.

    Signed-off-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Reviewed-by: Guennadi Liakhovetski <guennadi.liakhovetski@linux.intel.com>
```

```
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20211119192621.4096077-8-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 77755db2765d260201d102d58763ad0a9d99b962
Author: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
Date:   Fri Nov 19 21:26:17 2021 +0200

    ASoC: SOF: topology: remove sof_load_pipeline_ipc()

    commit 7cc7b9ba21d4978d19f0e3edc2b00d44c9d66ff6 upstream.

    Remove the function sof_load_pipeline_ipc() and directly
    send the IPC instead. The pipeline core is already enabled
    with the call to sof_pipeline_core_enable() in sof_widget_setup().

    Signed-off-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Reviewed-by: Guennadi Liakhovetski <guennadi.liakhovetski@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20211119192621.4096077-7-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 182251a7e992b98215d05c9e84460d7808773e06
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Mon Jan 24 14:05:34 2022 +0100

    KVM: selftests: Test KVM_SET_CPUID2 after KVM_RUN

    commit ecebb966acaab2466d9857d1cc435ee1fc9eee50 upstream.

    KVM forbids KVM_SET_CPUID2 after KVM_RUN was performed on a vCPU unless
    the supplied CPUID data is equal to what was previously set. Test this.

    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20220117150542.2176196-5-vkuznets@redhat.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 319fcc941e7ff09d3391978ed4a04b797615916d
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Mon Jan 24 14:05:33 2022 +0100

    KVM: selftests: Rename 'get_cpuid_test' to 'cpuid_test'

    commit 9e6d484f9991176269607bb3c54a494e32eab27a upstream.

    In preparation to reusing the existing 'get_cpuid_test' for testing
    "KVM_SET_CPUID{,2} after KVM_RUN" rename it to 'cpuid_test' to avoid
    the confusion.

    No functional change intended.

    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20220117150542.2176196-4-vkuznets@redhat.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 24e7590c60aa9487b8e43583dc9885f62f8216c1
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
```

```
Date:   Mon Jan 24 14:05:32 2022 +0100

    KVM: x86: Partially allow KVM_SET_CPUID{,2} after KVM_RUN

    commit c6617c61e8fe44b9e9fdfede921f61cac6b5149d upstream.

    Commit feb627e8d6f6 ("KVM: x86: Forbid KVM_SET_CPUID{,2} after KVM_RUN")
    forbade changing CPUID altogether but unfortunately this is not fully
    compatible with existing VMMs. In particular, QEMU reuses vCPU fds for
    CPU hotplug after unplug and it calls KVM_SET_CPUID2. Instead of full ban,
    check whether the supplied CPUID data is equal to what was previously set.

    Reported-by: Igor Mammedov <imammedo@redhat.com>
    Fixes: feb627e8d6f6 ("KVM: x86: Forbid KVM_SET_CPUID{,2} after KVM_RUN")
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20220117150542.2176196-3-vkuznets@redhat.com>
    Cc: stable@vger.kernel.org
    [Do not call kvm_find_cpuid_entry repeatedly. - Paolo]
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c34671d9f3b11b3eee75dacabf5c6e0d8a97cf8a
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Mon Jan 24 14:05:31 2022 +0100

    KVM: x86: Do runtime CPUID update before updating vcpu->arch.cpuid_entries

    commit ee3a5f9e3d9bf94159f3cc80da542fbe83502dd8 upstream.

    kvm_update_cpuid_runtime() mangles CPUID data coming from userspace
    VMM after updating 'vcpu->arch.cpuid_entries', this makes it
    impossible to compare an update with what was previously
    supplied. Introduce __kvm_update_cpuid_runtime() version which can be
    used to tweak the input before it goes to 'vcpu->arch.cpuid_entries'
    so the upcoming update check can compare tweaked data.

    No functional change intended.

    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20220117150542.2176196-2-vkuznets@redhat.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 32fb2bf4cdb550540d35839a0230f74051912af6
Author: Andrey Konovalov <andreyknvl@gmail.com>
Date:   Wed Jan 19 18:09:28 2022 -0800

    lib/test_meminit: destroy cache in kmem_cache_alloc_bulk() test

    commit e073e5ef90298d2d6e5e7f04b545a0815e92110c upstream.

    Make do_kmem_cache_size_bulk() destroy the cache it creates.

    Link:
https://lkml.kernel.org/r/aced20a94bf04159a139f0846e41d38a1537debb.1640018297.git.andreyknvl
@google.com
    Fixes: 03a9349ac0e0 ("lib/test_meminit: add a kmem_cache_alloc_bulk() test")
    Signed-off-by: Andrey Konovalov <andreyknvl@google.com>
    Reviewed-by: Marco Elver <elver@google.com>
    Cc: Alexander Potapenko <glider@google.com>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: Andrey Ryabinin <ryabinin.a.a@gmail.com>
    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 c5ed18a4ba1ba4579ab5539e9b187d3f0d5067e1
Author: Moshe Tal <moshet@nvidia.com>
Date:   Sun Jan 16 19:39:29 2022 +0200

    bonding: Fix extraction of ports from the packet headers

    commit 429e3d123d9a50cc9882402e40e0ac912d88cfcf upstream.

    Wrong hash sends single stream to multiple output interfaces.

    The offset calculation was relative to skb->head, fix it to be relative
    to skb->data.

    Fixes: a815bde56b15 ("net, bonding: Refactor bond_xmit_hash for use with
    xdp_buff")
    Reviewed-by: Jussi Maki <joamaki@gmail.com>
    Reviewed-by: Saeed Mahameed <saeedm@nvidia.com>
    Reviewed-by: Gal Pressman <gal@nvidia.com>
    Signed-off-by: Moshe Tal <moshet@nvidia.com>
    Acked-by: Jay Vosburgh <jay.vosburgh@canonical.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 017ef7c25c8394753144391dbd2c1667b98e31ab
Author: Alistair Popple <apopple@nvidia.com>
Date:   Fri Jan 14 14:09:31 2022 -0800

    mm/hmm.c: allow VM_MIXEDMAP to work with hmm_range_fault

    commit 87c01d57fa23de82fff593a7d070933d08755801 upstream.

    hmm_range_fault() can be used instead of get_user_pages() for devices
    which allow faulting however unlike get_user_pages() it will return an
    error when used on a VM_MIXEDMAP range.

    To make hmm_range_fault() more closely match get_user_pages() remove
    this restriction.  This requires dealing with the !ARCH_HAS_PTE_SPECIAL
    case in hmm_vma_handle_pte().  Rather than replicating the logic of
    vm_normal_page() call it directly and do a check for the zero pfn
    similar to what get_user_pages() currently does.

    Also add a test to hmm selftest to verify functionality.

    Link: https://lkml.kernel.org/r/20211104012001.2555676-1-apopple@nvidia.com
    Fixes: da4c3c735ea4 ("mm/hmm/mirror: helper to snapshot CPU page table")
    Signed-off-by: Alistair Popple <apopple@nvidia.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Cc: Jerome Glisse <jglisse@redhat.com>
    Cc: John Hubbard <jhubbard@nvidia.com>
    Cc: Zi Yan <ziy@nvidia.com>
    Cc: Ralph Campbell <rcampbell@nvidia.com>
    Cc: Felix Kuehling <Felix.Kuehling@amd.com>
    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 ad1ee1b1ff2f87c6042be83bd47fad1ead7834e7
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Jan 14 06:57:24 2022 +0000

    lib82596: Fix IRQ check in sni_82596_probe
```

```
    commit 99218cbf81bf21355a3de61cd46a706d36e900e6 upstream.

    platform_get_irq() returns negative error number instead 0 on failure.
    And the doc of platform_get_irq() provides a usage example:

        int irq = platform_get_irq(pdev, 0);
        if (irq < 0)
            return irq;

    Fix the check of return value to catch errors correctly.

    Fixes: 115978859272 ("i825xx: Move the Intel 82586/82593/82596 based drivers")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cf089f0713c0692bb90cb758d0a5cb971974fbf7
Author: Matthias Schiffer <matthias.schiffer@ew.tq-group.com>
Date:   Thu Jan 13 09:19:18 2022 +0100

    scripts/dtc: dtx_diff: remove broken example from help text

    commit d8adf5b92a9d2205620874d498c39923ecea8749 upstream.

    dtx_diff suggests to use <(...) syntax to pipe two inputs into it, but
    this has never worked: The /proc/self/fds/... paths passed by the shell
    will fail the `[ -f "${dtx}" ] && [ -r "${dtx}" ]` check in compile_to_dts,
    but even with this check removed, the function cannot work: hexdump will
    eat up the DTB magic, making the subsequent dtc call fail, as a pipe
    cannot be rewound.

    Simply remove this broken example, as there is already an alternative one
    that works fine.

    Fixes: 10eadc253ddf ("dtc: create tool to diff device trees")
    Signed-off-by: Matthias Schiffer <matthias.schiffer@ew.tq-group.com>
    Reviewed-by: Frank Rowand <frank.rowand@sony.com>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20220113081918.10387-1-matthias.schiffer@ew.tq-group.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2803127401d5f33cd7c999be9f4540443441d4b
Author: Maxim Mikityanskiy <maximmi@nvidia.com>
Date:   Wed Jan 12 12:28:05 2022 +0200

    sch_api: Don't skip qdisc attach on ingress

    commit de2d807b294d3d2ce5e59043ae2634016765d076 upstream.

    The attach callback of struct Qdisc_ops is used by only a few qdiscs:
    mq, mqprio and htb. qdisc_graft() contains the following logic
    (pseudocode):

        if (!qdisc->ops->attach) {
            if (ingress)
                do ingress stuff;
            else
                do egress stuff;
        }
        if (!ingress) {
            ...
            if (qdisc->ops->attach)
                qdisc->ops->attach(qdisc);
        } else {
            ...
```

```
        }

    As we see, the attach callback is not called if the qdisc is being
    attached to ingress (TC_H_INGRESS). That wasn't a problem for mq and
    mqprio, since they contain a check that they are attached to TC_H_ROOT,
    and they can't be attached to TC_H_INGRESS anyway.

    However, the commit cited below added the attach callback to htb. It is
    needed for the hardware offload, but in the non-offload mode it
    simulates the "do egress stuff" part of the pseudocode above. The
    problem is that when htb is attached to ingress, neither "do ingress
    stuff" nor attach() is called. It results in an inconsistency, and the
    following message is printed to dmesg:

    unregister_netdevice: waiting for lo to become free. Usage count = 2

    This commit addresses the issue by running "do ingress stuff" in the
    ingress flow even in the attach callback is present, which is fine,
    because attach isn't going to be called afterwards.

    The bug was found by syzbot and reported by Eric.

    Fixes: d03b195b5aa0 ("sch_htb: Hierarchical QoS hardware offload")
    Signed-off-by: Maxim Mikityanskiy <maximmi@nvidia.com>
    Reported-by: Eric Dumazet <edumazet@google.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b98752e28df1d27d945f03568e2ba51ee33bd55
Author: Sam Protsenko <semen.protsenko@linaro.org>
Date:   Sun Nov 21 18:56:36 2021 +0200

    dt-bindings: watchdog: Require samsung,syscon-phandle for Exynos7

    commit 33950f9a36aca55c2b1e6062d9b29f3e97f91c40 upstream.

    Exynos7 watchdog driver is clearly indicating that its dts node must
    define syscon phandle property. That was probably forgotten, so add it.

    Signed-off-by: Sam Protsenko <semen.protsenko@linaro.org>
    Fixes: 2b9366b66967 ("watchdog: s3c2410_wdt: Add support for Watchdog device on
Exynos7")
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Reviewed-by: Rob Herring <robh@kernel.org>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20211107202943.8859-2-semen.protsenko@linaro.org
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit adf5cc5429639eee0477c6f6773d19b79945658c
Author: Alexander Stein <alexander.stein@mailbox.org>
Date:   Sun Dec 19 10:41:55 2021 +0100

    dt-bindings: display: meson-vpu: Add missing amlogic,canvas property

    commit 640f35b871d29cd685ce0ea0762636381beeb98a upstream.

    This property was already mentioned in the old textual bindings
    amlogic,meson-vpu.txt, but got dropped during conversion.
    Adding it back similar to amlogic,gx-vdec.yaml.

    Fixes: 6b9ebf1e0e67 ("dt-bindings: display: amlogic, meson-vpu: convert to yaml")
    Signed-off-by: Alexander Stein <alexander.stein@mailbox.org>
```

```
    Acked-by: Rob Herring <robh@kernel.org>
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Reviewed-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211219094155.177206-1-
alexander.stein@mailbox.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7e35b8e5be91cb40c8367ce9d3e71fd39be3a371
Author: Alexander Stein <alexander.stein@mailbox.org>
Date:   Thu Dec 23 13:24:32 2021 +0100

    dt-bindings: display: meson-dw-hdmi: add missing sound-name-prefix property

    commit 22bf4047d26980807611b7e2030803db375afd87 upstream.

    This is used in meson-gx and meson-g12. Add the property to the binding.
    This fixes the dtschema warning:
    hdmi-tx@c883a000: 'sound-name-prefix' does not match any of the
    regexes: 'pinctrl-[0-9]+'

    Signed-off-by: Alexander Stein <alexander.stein@mailbox.org>
    Fixes: 376bf52deef5 ("dt-bindings: display: amlogic, meson-dw-hdmi: convert to yaml")
    Acked-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211223122434.39378-2-
alexander.stein@mailbox.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fb805465fae12195b1d22d805157c6d1eb5d4c3e
Author: Tom Rix <trix@redhat.com>
Date:   Tue Jan 18 05:41:10 2022 -0800

    net: mscc: ocelot: fix using match before it is set

    commit baa59504c1cd0cca7d41954a45ee0b3dc78e41a0 upstream.

    Clang static analysis reports this issue
    ocelot_flower.c:563:8: warning: 1st function call argument
      is an uninitialized value
        !is_zero_ether_addr(match.mask->dst)) {
        ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    The variable match is used before it is set.  So move the
    block.

    Fixes: 75944fda1dfe ("net: mscc: ocelot: offload ingress skbedit and vlan actions to
VCAP IS1")
    Signed-off-by: Tom Rix <trix@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e11efb9dc0f81416bb57e9f21974adf404332d78
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Tue Jan 18 13:08:12 2022 +0200

    net: phy: micrel: use kszphy_suspend()/kszphy_resume for irq aware devices

    commit f1131b9c23fb4a3540a774828ff49f421619f902 upstream.

    On a setup with KSZ9131 and MACB drivers it happens on suspend path, from
    time to time, that the PHY interrupt arrives after PHY and MACB were
    suspended (PHY via genphy_suspend(), MACB via macb_suspend()). In this
    case the phy_read() at the beginning of kszphy_handle_interrupt() will
    fail (as MACB driver is suspended at this time) leading to phy_error()
```

```
    being called and a stack trace being displayed on console. To solve this
    .suspend/.resume functions for all KSZ devices implementing
    .handle_interrupt were replaced with kszphy_suspend()/kszphy_resume()
    which disable/enable interrupt before/after calling
    genphy_suspend()/genphy_resume().

    The fix has been adapted for all KSZ devices which implements
    .handle_interrupt but it has been tested only on KSZ9131.

    Fixes: 59ca4e58b917 ("net: phy: micrel: implement generic .handle_interrupt() callback")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f0274b3dfc7140be153092f17371cddd011f5bd3
Author: Ard Biesheuvel <ardb@kernel.org>
Date:   Tue Jan 18 11:22:04 2022 +0100

    net: cpsw: avoid alignment faults by taking NET_IP_ALIGN into account

    commit 1771afd47430f5e95c9c3a2e3a8a63e67402d3fe upstream.

    Both versions of the CPSW driver declare a CPSW_HEADROOM_NA macro that
    takes NET_IP_ALIGN into account, but fail to use it appropriately when
    storing incoming packets in memory. This results in the IPv4 source and
    destination addresses to appear misaligned in memory, which causes
    aligment faults that need to be fixed up in software.

    So let's switch from CPSW_HEADROOM to CPSW_HEADROOM_NA where needed.
    This gets rid of any alignment faults on the RX path on a Beaglebone
    White.

    Fixes: 9ed4050c0d75 ("net: ethernet: ti: cpsw: add XDP support")
    Cc: Grygorii Strashko <grygorii.strashko@ti.com>
    Cc: Ilias Apalodimas <ilias.apalodimas@linaro.org>
    Signed-off-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2526941e6dad8b8c787cb973f28a6f6fa19957f
Author: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
Date:   Mon Jan 17 14:52:33 2022 +0000

    net: sfp: fix high power modules without diagnostic monitoring

    commit 5765cee119bf5a36c94d20eceb37c445508934be upstream.

    Commit 7cfa9c92d0a3 ("net: sfp: avoid power switch on address-change
    modules") unintetionally changed the semantics for high power modules
    without the digital diagnostics monitoring. We repeatedly attempt to
    read the power status from the non-existing 0xa2 address in a futile
    hope this failure is temporary:

    [    8.856051] sfp sfp-eth3: module NTT              0000000000000000 rev 0000  sn
0000000000000000 dc 160408
    [    8.865843] mvpp2 f4000000.ethernet eth3: switched to inband/1000base-x link mode
    [    8.873469] sfp sfp-eth3: Failed to read EEPROM: -5
    [    8.983251] sfp sfp-eth3: Failed to read EEPROM: -5
    [    9.103250] sfp sfp-eth3: Failed to read EEPROM: -5

    We previosuly assumed such modules were powered up in the correct mode,
    continuing without further configuration as long as the required power
    class was supported by the host.
```

```
    Restore this behaviour, while preserving the intent of subsequent
    patches to avoid the "Address Change Sequence not supported" warning
    if we are not going to be accessing the DDM address.

    Fixes: 7cfa9c92d0a3 ("net: sfp: avoid power switch on address-change modules")
    Reported-by: 照山周一郎 <teruyama@springboard-inc.jp>
    Tested-by: 照山周一郎 <teruyama@springboard-inc.jp>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 32eccf163a2ff51abb04c1c01d27c84b5d5083c8
Author: Horatiu Vultur <horatiu.vultur@microchip.com>
Date:   Mon Jan 17 13:53:00 2022 +0100

    net: ocelot: Fix the call to switchdev_bridge_port_offload

    commit c0b7f7d7e0ad44f35745c01964b3fa2833e298cb upstream.

    In the blamed commit, the call to the function
    switchdev_bridge_port_offload was passing the wrong argument for
    atomic_nb. It was ocelot_netdevice_nb instead of ocelot_swtchdev_nb.
    This patch fixes this issue.

    Fixes: 4e51bf44a03af6 ("net: bridge: move the switchdev object replay helpers to "push"
mode")
    Signed-off-by: Horatiu Vultur <horatiu.vultur@microchip.com>
    Reviewed-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 843300ae55303ddfa1fd874618d3f0186901aeeb
Author: Tom Rix <trix@redhat.com>
Date:   Sat Jan 15 09:49:18 2022 -0800

    net: ethernet: mtk_eth_soc: fix error checking in mtk_mac_config()

    commit 214b3369ab9b0a6f28d6c970220c209417edbc65 upstream.

    Clang static analysis reports this problem
    mtk_eth_soc.c:394:7: warning: Branch condition evaluates
      to a garbage value
                  if (err)
                      ^~~

    err is not initialized and only conditionally set.
    So intitialize err.

    Fixes: 7e538372694b ("net: ethernet: mediatek: Re-add support SGMII")
    Signed-off-by: Tom Rix <trix@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 09f65e277cc812f4d37af478da59c2a900cc7fed
Author: Slark Xiao <slark_xiao@163.com>
Date:   Sat Jan 15 10:34:30 2022 +0800

    net: wwan: Fix MRU mismatch issue which may lead to data connection lost

    commit f542cdfa3083a309e3caafbbdf41490c4935492a upstream.

    In pci_generic.c there is a 'mru_default' in struct mhi_pci_dev_info.
    This value shall be used for whole mhi if it's given a value for a specific product.
    But in function mhi_net_rx_refill_work(), it's still using hard code value
MHI_DEFAULT_MRU.
```

```
    'mru_default' shall have higher priority than MHI_DEFAULT_MRU.
    And after checking, this change could help fix a data connection lost issue.

    Fixes: 5c2c85315948 ("bus: mhi: pci-generic: configurable network interface MRU")
    Signed-off-by: Shujun Wang <wsj20369@163.com>
    Signed-off-by: Slark Xiao <slark_xiao@163.com>
    Reviewed-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7356267af5bf284cde6e78fe5d81d06a917f19af
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Fri Jan 14 15:36:37 2022 +0200

    net: mscc: ocelot: don't dereference NULL pointers with shared tc filters

    commit 80f15f3bef9e9c2cc29888a6773df44de0a0c65f upstream.

    The following command sequence:

    tc qdisc del dev swp0 clsact
    tc qdisc add dev swp0 ingress_block 1 clsact
    tc qdisc add dev swp1 ingress_block 1 clsact
    tc filter add block 1 flower action drop
    tc qdisc del dev swp0 clsact

    produces the following NPD:

    Unable to handle kernel NULL pointer dereference at virtual address 0000000000000014
    pc : vcap_entry_set+0x14/0x70
    lr : ocelot_vcap_filter_del+0x198/0x234
    Call trace:
     vcap_entry_set+0x14/0x70
     ocelot_vcap_filter_del+0x198/0x234
     ocelot_cls_flower_destroy+0x94/0xe4
     felix_cls_flower_del+0x70/0x84
     dsa_slave_setup_tc_block_cb+0x13c/0x60c
     dsa_slave_setup_tc_block_cb_ig+0x20/0x30
     tc_setup_cb_reoffload+0x44/0x120
     fl_reoffload+0x280/0x320
     tcf_block_playback_offloads+0x6c/0x184
     tcf_block_unbind+0x80/0xe0
     tcf_block_setup+0x174/0x214
     tcf_block_offload_cmd.isra.0+0x100/0x13c
     tcf_block_offload_unbind+0x5c/0xa0
     __tcf_block_put+0x54/0x174
     tcf_block_put_ext+0x5c/0x74
     clsact_destroy+0x40/0x60
     qdisc_destroy+0x4c/0x150
     qdisc_put+0x70/0x90
     qdisc_graft+0x3f0/0x4c0
     tc_get_qdisc+0x1cc/0x364
     rtnetlink_rcv_msg+0x124/0x340

    The reason is that the driver isn't prepared to receive two tc filters
    with the same cookie. It unconditionally creates a new struct
    ocelot_vcap_filter for each tc filter, and it adds all filters with the
    same identifier (cookie) to the ocelot_vcap_block.

    The problem is here, in ocelot_vcap_filter_del():

            /* Gets index of the filter */
            index = ocelot_vcap_block_get_filter_index(block, filter);
            if (index < 0)
                    return index;
```

```
		/* Delete filter */
		ocelot_vcap_block_remove_filter(ocelot, block, filter);

		/* Move up all the blocks over the deleted filter */
		for (i = index; i < block->count; i++) {
			struct ocelot_vcap_filter *tmp;

			tmp = ocelot_vcap_block_find_filter_by_index(block, i);
			vcap_entry_set(ocelot, i, tmp);
		}
```

what will happen is ocelot_vcap_block_get_filter_index() will return the index (@index) of the first filter found with that cookie. This is _not_ the index of _this_ filter, but the other one with the same cookie, because ocelot_vcap_filter_equal() gets fooled.

Then later, ocelot_vcap_block_remove_filter() is coded to remove all filters that are ocelot_vcap_filter_equal() with the passed @filter. So unexpectedly, both filters get deleted from the list.

Then ocelot_vcap_filter_del() will attempt to move all the other filters up, again finding them by index (@i). The block count is 2, @index was 0, so it will attempt to move up filter @i=0 and @i=1. It assigns tmp = ocelot_vcap_block_find_filter_by_index(block, i), which is now a NULL pointer because ocelot_vcap_block_remove_filter() has removed more than one filter.

As far as I can see, this problem has been there since the introduction of tc offload support, however I cannot test beyond the blamed commit due to hardware availability. In any case, any fix cannot be backported that far, due to lots of changes to the code base.

Therefore, let's go for the correct solution, which is to not call ocelot_vcap_filter_add() and ocelot_vcap_filter_del(), unless the filter is actually unique and not shared. For the shared filters, we should just modify the ingress port mask and call ocelot_vcap_filter_replace(), a function introduced by commit 95706be13b9f ("net: mscc: ocelot: create a function that replaces an existing VCAP filter"). This way, block->rules will only contain filters with unique cookies, by design.

Fixes: 07d985eef073 ("net: dsa: felix: Wire up the ocelot cls_flower methods")
Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1580365cf1a091dd5382cb350e51a14df56d97ec
Author: Sergey Shtylyov <s.shtylyov@omp.ru>
Date: Thu Jan 13 22:46:07 2022 +0300

bcmgenet: add WOL IRQ check

commit 9deb48b53e7f4056c2eaa2dc2ee3338df619e4f6 upstream.

The driver neglects to check the result of platform_get_irq_optional()'s call and blithely passes the negative error codes to devm_request_irq() (which takes *unsigned* IRQ #), causing it to fail with -EINVAL. Stop calling devm_request_irq() with the invalid IRQ #s.

Fixes: 8562056f267d ("net: bcmgenet: request Wake-on-LAN interrupt")
Signed-off-by: Sergey Shtylyov <s.shtylyov@omp.ru>
Acked-by: Florian Fainelli <f.fainelli@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

```
commit 79ec330d03983095ef29c11d07943700fb878d31
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Wed Jan 12 22:21:27 2022 +0200

    net: mscc: ocelot: don't let phylink re-enable TX PAUSE on the NPI port

    commit 33cb0ff30cff104e753f7882c99e54cf67ea7903 upstream.

    Since commit b39648079db4 ("net: mscc: ocelot: disable flow control on
    NPI interface"), flow control should be disabled on the DSA CPU port
    when used in NPI mode.

    However, the commit blamed in the Fixes: tag below broke this, because
    it allowed felix_phylink_mac_link_up() to overwrite SYS_PAUSE_CFG_PAUSE_ENA
    for the DSA CPU port.

    This issue became noticeable since the device tree update from commit
    8fcea7be5736 ("arm64: dts: ls1028a: mark internal links between Felix
    and ENETC as capable of flow control").

    The solution is to check whether this is the currently configured NPI
    port from ocelot_phylink_mac_link_up(), and to not modify the statically
    disabled PAUSE frame transmission if it is.

    When the port is configured for lossless mode as opposed to tail drop
    mode, but the link partner (DSA master) doesn't observe the transmitted
    PAUSE frames, the switch termination throughput is much worse, as can be
    seen below.

    Before:

    root@debian:~# iperf3 -c 192.168.100.2
    Connecting to host 192.168.100.2, port 5201
    [  5] local 192.168.100.1 port 37504 connected to 192.168.100.2 port 5201
    [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
    [  5]   0.00-1.00   sec  28.4 MBytes   238 Mbits/sec  357   22.6 KBytes
    [  5]   1.00-2.00   sec  33.6 MBytes   282 Mbits/sec  426   19.8 KBytes
    [  5]   2.00-3.00   sec  34.0 MBytes   285 Mbits/sec  343   21.2 KBytes
    [  5]   3.00-4.00   sec  32.9 MBytes   276 Mbits/sec  354   22.6 KBytes
    [  5]   4.00-5.00   sec  32.3 MBytes   271 Mbits/sec  297   18.4 KBytes
    ^C[  5]   5.00-5.06   sec  2.05 MBytes   270 Mbits/sec   45   19.8 KBytes
    - - - - - - - - - - - - - - - - - - - - - - - - -
    [ ID] Interval           Transfer     Bitrate         Retr
    [  5]   0.00-5.06   sec   163 MBytes   271 Mbits/sec  1822             sender
    [  5]   0.00-5.06   sec  0.00 Bytes  0.00 bits/sec                  receiver

    After:

    root@debian:~# iperf3 -c 192.168.100.2
    Connecting to host 192.168.100.2, port 5201
    [  5] local 192.168.100.1 port 49470 connected to 192.168.100.2 port 5201
    [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
    [  5]   0.00-1.00   sec   112 MBytes   941 Mbits/sec  259    143 KBytes
    [  5]   1.00-2.00   sec   110 MBytes   920 Mbits/sec  329    144 KBytes
    [  5]   2.00-3.00   sec   112 MBytes   936 Mbits/sec  255    144 KBytes
    [  5]   3.00-4.00   sec   110 MBytes   927 Mbits/sec  355    105 KBytes
    [  5]   4.00-5.00   sec   110 MBytes   926 Mbits/sec  350    156 KBytes
    [  5]   5.00-6.00   sec   110 MBytes   925 Mbits/sec  305    148 KBytes
    [  5]   6.00-7.00   sec   110 MBytes   924 Mbits/sec  320    143 KBytes
    [  5]   7.00-8.00   sec   110 MBytes   925 Mbits/sec  273   97.6 KBytes
    [  5]   8.00-9.00   sec   109 MBytes   913 Mbits/sec  299    141 KBytes
    [  5]   9.00-10.00  sec   110 MBytes   922 Mbits/sec  287    146 KBytes
    - - - - - - - - - - - - - - - - - - - - - - - - -
    [ ID] Interval           Transfer     Bitrate         Retr
    [  5]   0.00-10.00  sec  1.08 GBytes   926 Mbits/sec  3032             sender
```

```
    [  5]   0.00-10.00  sec  1.08 GBytes   925 Mbits/sec                  receiver

    Fixes: de274be32cb2 ("net: dsa: felix: set TX flow control according to the
phylink_mac_link_up resolution")
    Reported-by: Xiaoliang Yang <xiaoliang.yang_1@nxp.com>
    Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Reviewed-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit acac3b44c43570b03ee2ea212e6bb093e4dd3f27
Author: Kevin Bracey <kevin@bracey.fi>
Date:   Wed Jan 12 19:02:10 2022 +0200

    net_sched: restore "mpu xxx" handling

    commit fb80445c438c78b40b547d12b8d56596ce4ccfeb upstream.

    commit 56b765b79e9a ("htb: improved accuracy at high rates") broke
    "overhead X", "linklayer atm" and "mpu X" attributes.

    "overhead X" and "linklayer atm" have already been fixed. This restores
    the "mpu X" handling, as might be used by DOCSIS or Ethernet shaping:

        tc class add ... htb rate X overhead 4 mpu 64

    The code being fixed is used by htb, tbf and act_police. Cake has its
    own mpu handling. qdisc_calculate_pkt_len still uses the size table
    containing values adjusted for mpu by user space.

    iproute2 tc has always passed mpu into the kernel via a tc_ratespec
    structure, but the kernel never directly acted on it, merely stored it
    so that it could be read back by `tc class show`.

    Rather, tc would generate length-to-time tables that included the mpu
    (and linklayer) in their construction, and the kernel used those tables.

    Since v3.7, the tables were no longer used. Along with "mpu", this also
    broke "overhead" and "linklayer" which were fixed in 01cb71d2d47b
    ("net_sched: restore "overhead xxx" handling", v3.10) and 8a8e3d84b171
    ("net_sched: restore "linklayer atm" handling", v3.11).

    "overhead" was fixed by simply restoring use of tc_ratespec::overhead -
    this had originally been used by the kernel but was initially omitted
    from the new non-table-based calculations.

    "linklayer" had been handled in the table like "mpu", but the mode was
    not originally passed in tc_ratespec. The new implementation was made to
    handle it by getting new versions of tc to pass the mode in an extended
    tc_ratespec, and for older versions of tc the table contents were analysed
    at load time to deduce linklayer.

    As "mpu" has always been given to the kernel in tc_ratespec,
    accompanying the mpu-based table, we can restore system functionality
    with no userspace change by making the kernel act on the tc_ratespec
    value.

    Fixes: 56b765b79e9a ("htb: improved accuracy at high rates")
    Signed-off-by: Kevin Bracey <kevin@bracey.fi>
    Cc: Eric Dumazet <edumazet@google.com>
    Cc: Jiri Pirko <jiri@resnulli.us>
    Cc: Vimalkumar <j.vimal@gmail.com>
    Link: https://lore.kernel.org/r/20220112170210.1014351-1-kevin@bracey.fi
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 6279192add94a4d01589186d5a936bdfbc7d9ee7
Author: Alex Elder <elder@linaro.org>
Date:   Wed Jan 12 07:30:10 2022 -0600

    net: ipa: fix atomic update in ipa_endpoint_replenish()

    commit 6c0e3b5ce94947b311348c367db9e11dcb2ccc93 upstream.

    In ipa_endpoint_replenish(), if an error occurs when attempting to
    replenish a receive buffer, we just quit and try again later.  In
    that case we increment the backlog count to reflect that the attempt
    was unsuccessful.  Then, if the add_one flag was true we increment
    the backlog again.

    This second increment is not included in the backlog local variable
    though, and its value determines whether delayed work should be
    scheduled.  This is a bug.

    Fix this by determining whether 1 or 2 should be added to the
    backlog before adding it in a atomic_add_return() call.

    Reviewed-by: Matthias Kaehlcke <mka@chromium.org>
    Fixes: 84f9bd12d46db ("soc: qcom: ipa: IPA endpoints")
    Signed-off-by: Alex Elder <elder@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c1f1691ef84fa6d38fa5e5148eca073145e97ffa
Author: Jie Wang <wangjie125@huawei.com>
Date:   Wed Jan 12 20:54:18 2022 +0800

    net: bonding: fix bond_xmit_broadcast return value error bug

    commit 4e5bd03ae34652cd932ab4c91c71c511793df75c upstream.

    In Linux bonding scenario, one packet is copied to several copies and sent
    by all slave device of bond0 in mode 3(broadcast mode). The mode 3 xmit
    function bond_xmit_broadcast() only ueses the last slave device's tx result
    as the final result. In this case, if the last slave device is down, then
    it always return NET_XMIT_DROP, even though the other slave devices xmit
    success. It may cause the tx statistics error, and cause the application
    (e.g. scp) consider the network is unreachable.

    For example, use the following command to configure server A.

    echo 3 > /sys/class/net/bond0/bonding/mode
    ifconfig bond0 up
    ifenslave bond0 eth0 eth1
    ifconfig bond0 192.168.1.125
    ifconfig eth0 up
    ifconfig eth1 down
    The slave device eth0 and eth1 are connected to server B(192.168.1.107).
    Run the ping 192.168.1.107 -c 3 -i 0.2 command, the following information
    is displayed.

    PING 192.168.1.107 (192.168.1.107) 56(84) bytes of data.
    64 bytes from 192.168.1.107: icmp_seq=1 ttl=64 time=0.077 ms
    64 bytes from 192.168.1.107: icmp_seq=2 ttl=64 time=0.056 ms
    64 bytes from 192.168.1.107: icmp_seq=3 ttl=64 time=0.051 ms

     192.168.1.107 ping statistics
    0 packets transmitted, 3 received

    Actually, the slave device eth0 of the bond successfully sends three
```

```
    ICMP packets, but the result shows that 0 packets are transmitted.

    Also if we use scp command to get remote files, the command end with the
    following printings.

    ssh_exchange_identification: read: Connection timed out

    So this patch modifies the bond_xmit_broadcast to return NET_XMIT_SUCCESS
    if one slave device in the bond sends packets successfully. If all slave
    devices send packets fail, the discarded packets stats is increased. The
    skb is released when there is no slave device in the bond or the last slave
    device is down.

    Fixes: ae46f184bc1f ("bonding: propagate transmit status")
    Signed-off-by: Jie Wang <wangjie125@huawei.com>
    Signed-off-by: Guangbin Huang <huangguangbin2@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bd0f4a3eea2c856507fdd9003bcbe826affe3007
Author: Miroslav Lichvar <mlichvar@redhat.com>
Date:   Tue Jan 11 16:10:53 2022 +0100

    net: fix sock_timestamping_bind_phc() to release device

    commit 2a4d75bfe41232608f5596a6d1369f92ccb20817 upstream.

    Don't forget to release the device in sock_timestamping_bind_phc() after
    it was used to get the vclock indices.

    Fixes: d463126e23f1 ("net: sock: extend SO_TIMESTAMPING for PHC binding")
    Signed-off-by: Miroslav Lichvar <mlichvar@redhat.com>
    Cc: Yangbo Lu <yangbo.lu@nxp.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7d8df20e4c74a6ed429395251acb6c84e61737d0
Author: David Heidelberg <david@ixit.cz>
Date:   Sat Oct 30 12:04:12 2021 +0200

    arm64: dts: qcom: msm8996: drop not documented adreno properties

    commit c41910f257a22dc406c60d8826b4a3b5398003a3 upstream.

    These properties aren't documented nor implemented in the driver.
    Drop them.

    Fixes warnings as:
    $ make dtbs_check DT_SCHEMA_FILES=Documentation/devicetree/bindings/display/msm/gpu.yaml
    ...
    arch/arm64/boot/dts/qcom/msm8996-mtp.dt.yaml: gpu@b00000: 'qcom,gpu-quirk-fault-detect-
mask', 'qcom,gpu-quirk-two-pass-use-wfi' do not match any of the regexes: 'pinctrl-[0-9]+'
            From schema: Documentation/devicetree/bindings/display/msm/gpu.yaml
    ...

    Fixes: 69cc3114ab0f ("arm64: dts: Add Adreno GPU definitions")
    Signed-off-by: David Heidelberg <david@ixit.cz>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211030100413.28370-1-david@ixit.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1d3ba3a9f4c060af25c12fb89f2455b47a476368
Author: Leon Romanovsky <leon@kernel.org>
Date:   Sun Nov 28 14:14:46 2021 +0200
```

```
    devlink: Remove misleading internal_flags from health reporter dump

    commit e9538f8270db24d272659e15841854c7ea11119e upstream.

    DEVLINK_CMD_HEALTH_REPORTER_DUMP_GET command doesn't have .doit callback
    and has no use in internal_flags at all. Remove this misleading assignment.

    Fixes: e44ef4e4516c ("devlink: Hang reporter's dump method on a dumpit cb")
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 14c6618913fc9670e1e2400d792f8830d2d6eff0
Author: Ian Rogers <irogers@google.com>
Date:   Fri Jan 14 22:28:52 2022 -0800

    perf metric: Fix metric_leader

    commit d3e2bb4359f70c8b1d09a6f8e2f57240aab0da3f upstream.

    Multiple events may have a metric_leader to aggregate into.

    This happens for uncore events where, for example, uncore_imc is
    expanded into uncore_imc_0, uncore_imc_1, etc.

    Such events all have the same metric_id and should aggregate into the
    first event.

    The change introducing metric_ids had a bug where the metric_id was
    compared to itself, creating an always true condition.

    Correct this by comparing the event in the metric_evlist and the
    metric_leader.

    Fixes: ec5c5b3d2c21b3f3 ("perf metric: Encode and use metric-id as qualifier")
    Signed-off-by: Ian Rogers <irogers@google.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Andi Kleen <ak@linux.intel.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: John Garry <john.garry@huawei.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Stephane Eranian <eranian@google.com>
    Link: http://lore.kernel.org/lkml/20220115062852.1959424-1-irogers@google.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3414cc48b53eba8d0d78ec772ff8332eeef0ccd2
Author: Zechuan Chen <chenzechuan1@huawei.com>
Date:   Tue Dec 28 19:13:38 2021 +0800

    perf probe: Fix ppc64 'perf probe add events failed' case

    commit 4624f199327a704dd1069aca1c3cadb8f2a28c6f upstream.

    Because of commit bf794bf52a80c627 ("powerpc/kprobes: Fix kallsyms
    lookup across powerpc ABIv1 and ABIv2"), in ppc64 ABIv1, our perf
    command eliminates the need to use the prefix "." at the symbol name.

    But when the command "perf probe -a schedule" is executed on ppc64
    ABIv1, it obtains two symbol address information through /proc/kallsyms,
    for example:

      cat /proc/kallsyms | grep -w schedule
```

```
       c000000000657020 T .schedule
       c000000000d4fdb8 D schedule

    The symbol "D schedule" is not a function symbol, and perf will print:
    "p:probe/schedule _text+13958584"Failed to write event: Invalid argument

    Therefore, when searching symbols from map and adding probe point for
    them, a symbol type check is added. If the type of symbol is not a
    function, skip it.

    Fixes: bf794bf52a80c627 ("powerpc/kprobes: Fix kallsyms lookup across powerpc ABIv1 and
ABIv2")
    Signed-off-by: Zechuan Chen <chenzechuan1@huawei.com>
    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Jianlin Lv <Jianlin.Lv@arm.com>
    Cc: Jin Yao <yao.jin@linux.intel.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Michael Ellerman <mpe@ellerman.id.au>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Naveen N. Rao <naveen.n.rao@linux.vnet.ibm.com>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Ravi Bangoria <ravi.bangoria@linux.ibm.com>
    Cc: Yang Jihong <yangjihong1@huawei.com>
    Link: https://lore.kernel.org/r/20211228111338.218602-1-chenzechuan1@huawei.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8f2530665fc18f96433de0f704cfa61d16c30e86
Author: Ian Rogers <irogers@google.com>
Date:   Thu Dec 23 10:39:47 2021 -0800

    perf test: Enable system wide for metricgroups test

    commit 0046686da0ef692a6381260c3aa44291187eafc9 upstream.

    Uncore events as group leaders fail in per-thread mode causing exit
    errors. Enable system-wide for metricgroup testing. This fixes the HPC
    metric group when tested on skylakex.

    Fixes: 4a87dea9e60fe100 ("perf test: Workload test of metric and metricgroups")
    Signed-off-by: Ian Rogers <irogers@google.com>
    Tested-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Andi Kleen <ak@linux.intel.com>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Cc: Stephane Eranian <eranian@google.com>
    Link: https://lore.kernel.org/r/20211223183948.3423989-1-irogers@google.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit af8bc222dbc1233e0b54936e8d469db144a8b89a
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Wed Dec 8 18:11:13 2021 +0100

    perf metricgroup: Fix use after free in metric__new()

    commit e000ea0beffb5497425054b151369fe37a792ece upstream.
```

```
    We shouldn't free() something that will be used in the next line, fix
    it.

    Fixes: b85a4d61d3022608 ("perf metric: Allow modifiers on metrics")
    Addresses-Coverity-ID: 1494000
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Andi Kleen <ak@linux.intel.com>
    Cc: Ian Rogers <irogers@google.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: John Garry <john.garry@huawei.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Link: http://lore.kernel.org/lkml/20211208171113.22089-1-jose.exposito89@gmail.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d719d6b0983cc7879d895777165acd7035b693df
Author: Uwe Kleine-König <uwe@kleine-koenig.org>
Date:   Fri Dec 3 22:05:44 2021 +0100

    perf tools: Drop requirement for libstdc++.so for libopencsd check

    commit ed17b1914978eddb2b01f2d34577f1c82518c650 upstream.

    It's possible to link against libopencsd_c_api without having
    libstdc++.so available, only libstdc++.so.6.0.28 (or whatever version is
    in use) needs to be available. The same holds true for libopencsd.so.
    When -lstdc++ (or -lopencsd) is explicitly passed to the linker however
    the .so file must be available.

    So wrap adding the dependencies into a check for static linking that
    actually requires adding them all. The same construct is already used
    for some other tests in the same file to reduce dependencies in the
    dynamic linking case.

    Fixes: 573cf5c9a152 ("perf build: Add missing -lstdc++ when linking with libopencsd")
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Uwe Kleine-König <uwe@kleine-koenig.org>
    Cc: Adrian Bunk <bunk@debian.org>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Branislav Rankov <branislav.rankov@arm.com>
    Cc: Diederik de Haas <didi.debian@cknow.org>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Link: https://lore.kernel.org/all/20211203210544.1137935-1-uwe@kleine-koenig.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 027e3c50fd6add84c45fd3b216a002ba4ec8e411
Author: Thomas Richter <tmricht@linux.ibm.com>
Date:   Wed Nov 24 10:03:43 2021 +0100

    perf cputopo: Fix CPU topology reading on s/390

    commit a6e62743621ea29bea461774c0bcc68e5de59068 upstream.

    Commit fdf1e29b6118c18f ("perf expr: Add metric literals for topology.")
    fails on s390:

     # ./perf test -Fv 7
       ...
```

```
      # FAILED tests/expr.c:173 #num_dies >= #num_packages
        ---- end ----
        Simple expression parser: FAILED!
      #
```

Investigating this issue leads to these functions:

```
      build_cpu_topology()
        +--> has_die_topology(void)
             {
                struct utsname uts;

                if (uname(&uts) < 0)
                       return false;
                if (strncmp(uts.machine, "x86_64", 6))
                       return false;
                ....
             }
```

which always returns false on s390. The caller build_cpu_topology()
checks has_die_topology() return value. On false the
the struct cpu_topology::die_cpu_list is not contructed and has zero
entries. This leads to the failing comparison: #num_dies >= #num_packages.
s390 of course has a positive number of packages.

Fix this by adding s390 architecture to support CPU die list.

Output after:

```
      # ./perf test -Fv 7
        7: Simple expression parser                                     :
        --- start ---
        division by zero
        syntax error
        ---- end ----
        Simple expression parser: Ok
      #
```

Fixes: fdf1e29b6118c18f ("perf expr: Add metric literals for topology.")
Reviewed-by: Ian Rogers <irogers@google.com>
Signed-off-by: Thomas Richter <tmricht@linux.ibm.com>
Cc: Heiko Carstens <hca@linux.ibm.com>
Cc: Ian Rogers <irogers@google.com>
Cc: Sumanth Korikkar <sumanthk@linux.ibm.com>
Cc: Sven Schnelle <svens@linux.ibm.com>
Cc: Vasily Gorbik <gor@linux.ibm.com>
Link: https://lore.kernel.org/r/20211124090343.9436-1-tmricht@linux.ibm.com
Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e2d4094004804da378795ad5d14f38c53db7a68
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date: Wed Dec 15 13:01:13 2021 +0200

    dmaengine: at_xdmac: Fix at_xdmac_lld struct definition

    commit 912f7c6f7fac273f40e621447cf17d14b50d6e5b upstream.

    The hardware channel next descriptor view structure contains just
    fields of 32 bits, while dma_addr_t can be of type u64 or u32
    depending on CONFIG_ARCH_DMA_ADDR_T_64BIT. Force u32 to comply with
    what the hardware expects.

    Fixes: e1f7c9eee707 ("dmaengine: at_xdmac: creation of the atmel eXtended DMA Controller
driver")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-11-tudor.ambarus@microchip.com


```
```

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

commit 910969e4db23188fc2dbac02e6b7c17359fa605f
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Dec 15 13:01:12 2021 +0200

    dmaengine: at_xdmac: Fix lld view setting

    commit 1385eb4d14d447cc5d744bc2ac34f43be66c9963 upstream.

    AT_XDMAC_CNDC_NDVIEW_NDV3 was set even for AT_XDMAC_MBR_UBC_NDV2,
    because of the wrong bit handling. Fix it.

    Fixes: ee0fe35c8dcd ("dmaengine: xdmac: Handle descriptor's view 3 registers")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-10-tudor.ambarus@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3e084113cbfcfbc7534a0e35e160ef684ecc3f30
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Dec 15 13:01:10 2021 +0200

    dmaengine: at_xdmac: Fix concurrency over xfers_list

    commit 18deddea9184b62941395889ff7659529c877326 upstream.

    Since tx_submit can be called from a hard IRQ, xfers_list must be
    protected with a lock to avoid concurency on the list's elements.
    Since at_xdmac_handle_cyclic() is called from a tasklet, spin_lock_irq
    is enough to protect from a hard IRQ.

    Fixes: e1f7c9eee707 ("dmaengine: at_xdmac: creation of the atmel eXtended DMA Controller
driver")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-8-tudor.ambarus@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9eb7797f34c210bc1e9649ac616c52380bd506c9
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Dec 15 13:01:06 2021 +0200

    dmaengine: at_xdmac: Print debug message after realeasing the lock

    commit 5edc24ac876a928f36f407a0fcdb33b94a3a210f upstream.

    It is desirable to do the prints without the lock held if possible, so
    move the print after the lock is released.

    Fixes: e1f7c9eee707 ("dmaengine: at_xdmac: creation of the atmel eXtended DMA Controller
driver")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-4-tudor.ambarus@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 81ad2352fb3d305a6092e095e3f6b6d474102316
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Dec 15 13:01:05 2021 +0200

    dmaengine: at_xdmac: Start transfer for cyclic channels in issue_pending

    commit e6af9b05bec63cd4d1de2a33968cd0be2a91282a upstream.
```

```
    Cyclic channels must too call issue_pending in order to start a transfer.
    Start the transfer in issue_pending regardless of the type of channel.
    This wrongly worked before, because in the past the transfer was started
    at tx_submit level when only a desc in the transfer list.

    Fixes: e1f7c9eee707 ("dmaengine: at_xdmac: creation of the atmel eXtended DMA Controller
driver")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-3-tudor.ambarus@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 73587fb39cb64de7a30fd4ce1e787172b401cb6a
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Dec 15 13:01:04 2021 +0200

    dmaengine: at_xdmac: Don't start transactions at tx_submit level

    commit bccfb96b59179d4f96cbbd1ddff8fac6d335eae4 upstream.

    tx_submit is supposed to push the current transaction descriptor to a
    pending queue, waiting for issue_pending() to be called. issue_pending()
    must start the transfer, not tx_submit(), thus remove
    at_xdmac_start_xfer() from at_xdmac_tx_submit(). Clients of at_xdmac that
    assume that tx_submit() starts the transfer must be updated and call
    dma_async_issue_pending() if they miss to call it (one example is
    atmel_serial).

    As the at_xdmac_start_xfer() is now called only from
    at_xdmac_advance_work() when !at_xdmac_chan_is_enabled(), the
    at_xdmac_chan_is_enabled() check is no longer needed in
    at_xdmac_start_xfer(), thus remove it.

    Fixes: e1f7c9eee707 ("dmaengine: at_xdmac: creation of the atmel eXtended DMA Controller
driver")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211215110115.191749-2-tudor.ambarus@microchip.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d63d9415a312fb601f6cae28f83a058e04b3882e
Author: Adrian Hunter <adrian.hunter@intel.com>
Date:   Wed Jan 12 10:50:57 2022 +0200

    perf script: Fix hex dump character output

    commit 62942e9fda9fd1def10ffcbd5e1c025b3c9eec17 upstream.

    Using grep -C with perf script -D can give erroneous results as grep loses
    lines due to non-printable characters, for example, below the 0020, 0060
    and 0070 lines are missing:

     $ perf script -D | grep -C10 AUX | head
     .  0010:  08 00 00 00 00 00 00 00 1f 00 00 00 00 00 00 00  ................
     .  0030:  01 00 00 00 00 00 00 00 00 04 00 00 00 00 00 00  ................
     .  0040:  00 08 00 00 00 00 00 00 02 00 00 00 00 00 00 00  ................
     .  0050:  00 00 00 00 00 00 00 00 01 00 00 00 00 00 00 00  ................
     .  0080:  02 00 00 00 00 00 00 00 1b 00 00 00 00 00 00 00  ................
     .  0090:  00 00 00 00 00 00 00 00                          ........

     0 0 0x450 [0x98]: PERF_RECORD_AUXTRACE_INFO type: 1
       PMU Type            8
       Time Shift          31
```

```
    perf's isprint() is a custom implementation from the kernel, but the
    kernel's _ctype appears to include characters from Latin-1 Supplement which
    is not compatible with, for example, UTF-8. Fix by checking also isascii().

    After:

     $ tools/perf/perf script -D | grep -C10 AUX | head
     .  0010:  08 00 00 00 00 00 00 00 1f 00 00 00 00 00 00 00  ................
     .  0020:  03 84 32 2f 00 00 00 00 63 7c 4f d2 fa ff ff ff  ..2/....c|O.....
     .  0030:  01 00 00 00 00 00 00 00 00 04 00 00 00 00 00 00  ................
     .  0040:  00 08 00 00 00 00 00 00 02 00 00 00 00 00 00 00  ................
     .  0050:  00 00 00 00 00 00 00 00 01 00 00 00 00 00 00 00  ................
     .  0060:  00 02 00 00 00 00 00 00 00 c0 03 00 00 00 00 00  ................
     .  0070:  e2 00 00 00 00 00 00 00 02 00 00 00 00 00 00 00  ................
     .  0080:  02 00 00 00 00 00 00 00 1b 00 00 00 00 00 00 00  ................
     .  0090:  00 00 00 00 00 00 00 00                          ........

    Fixes: 3052ba56bcb58904 ("tools perf: Move from sane_ctype.h obtained from git to the
Linux's original")
    Signed-off-by: Adrian Hunter <adrian.hunter@intel.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Link: http://lore.kernel.org/lkml/20220112085057.277205-1-adrian.hunter@intel.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b785349dd7878e7ed5e3b37f6d20b59a7e2882f0
Author: Guillaume Nault <gnault@redhat.com>
Date:   Mon Jan 10 14:43:11 2022 +0100

    libcxgb: Don't accidentally set RTO_ONLINK in cxgb_find_route()

    commit a915deaa9abe4fb3a440312c954253a6a733608e upstream.

    Mask the ECN bits before calling ip_route_output_ports(). The tos
    variable might be passed directly from an IPv4 header, so it may have
    the last ECN bit set. This interferes with the route lookup process as
    ip_route_output_key_hash() interpretes this bit specially (to restrict
    the route scope).

    Found by code inspection, compile tested only.

    Fixes: 804c2f3e36ef ("libcxgb,iw_cxgb4,cxgbit: add cxgb_find_route()")
    Signed-off-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3c71127fb5ecb4058cc39a3d97054a93cac185f7
Author: Guillaume Nault <gnault@redhat.com>
Date:   Mon Jan 10 14:43:09 2022 +0100

    gre: Don't accidentally set RTO_ONLINK in gre_fill_metadata_dst()

    commit f7716b318568b22fbf0e3be99279a979e217cf71 upstream.

    Mask the ECN bits before initialising ->flowi4_tos. The tunnel key may
    have the last ECN bit set, which will interfere with the route lookup
    process as ip_route_output_key_hash() interpretes this bit specially
    (to restrict the route scope).

    Found by code inspection, compile tested only.

    Fixes: 962924fa2b7a ("ip_gre: Refactor collect metatdata mode tunnel xmit to
ip_md_tunnel_xmit")
    Signed-off-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
```

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

commit 1e9c263b9b96e4d057ce7e2510bb0214a449b2ca
Author: Eli Cohen <elic@nvidia.com>
Date:   Wed Jan 5 13:46:41 2022 +0200

    vdpa/mlx5: Restore cur_num_vqs in case of failure in change_num_qps()

    commit 37e07e705888e4c3502f204e9c6785c9c2d6d86a upstream.

    Restore ndev->cur_num_vqs to the original value in case change_num_qps()
    fails.

    Fixes: 52893733f2c5 ("vdpa/mlx5: Add multiqueue support")
    Reviewed-by: Si-Wei Liu<si-wei.liu@oracle.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20220105114646.577224-10-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a23a9a627a2218cf6439e1baac3de8565606c5a2
Author: Eli Cohen <elic@nvidia.com>
Date:   Wed Jan 5 13:46:38 2022 +0200

    vdpa/mlx5: Fix config_attr_mask assignment

    commit e3137056e6dedee205fccd06da031a285c6e34f5 upstream.

    Fix VDPA_ATTR_DEV_NET_CFG_MACADDR assignment to be explicit 64 bit
    assignment.

    No issue was seen since the value is well below 64 bit max value.
    Nevertheless it needs to be fixed.

    Fixes: a007d940040c ("vdpa/mlx5: Support configuration of MAC")
    Reviewed-by: Si-Wei Liu <si-wei.liu@oracle.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20220105114646.577224-7-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6bd3710823e041aee128f27a80430ef473f41f11
Author: Guillaume Nault <gnault@redhat.com>
Date:   Mon Jan 10 14:43:06 2022 +0100

    xfrm: Don't accidentally set RTO_ONLINK in decode_session4()

    commit 23e7b1bfed61e301853b5e35472820d919498278 upstream.

    Similar to commit 94e2238969e8 ("xfrm4: strip ECN bits from tos field"),
    clear the ECN bits from iph->tos when setting ->flowi4_tos.
    This ensures that the last bit of ->flowi4_tos is cleared, so
    ip_route_output_key_hash() isn't going to restrict the scope of the
    route lookup.

    Use ~INET_ECN_MASK instead of IPTOS_RT_MASK, because we have no reason
    to clear the high order bits.

    Found by code inspection, compile tested only.

    Fixes: 4da3089f2b58 ("[IPSEC]: Use TOS when doing tunnel lookups")
    Signed-off-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
```

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

commit 5cd640f6060a35ead0d61fe5538b01829a2dfe6f
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Fri Dec 10 09:06:11 2021 +0200

    iwlwifi: fix Bz NMI behaviour

    commit fdfde0cb79264f88992e72b5a056a3a3284fcaad upstream.

    Contrary to what was stated before, the hardware hasn't changed
    the bits here yet. In any case, the new CSR is also directly
    (lower 16 bits) connected to UREG_DOORBELL_TO_ISR6, so if it
    still changes the changes would be there. Adjust the code and
    comments accordingly.

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Fixes: 6c0795f1a524 ("iwlwifi: implement Bz NMI behaviour")
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211210090244.75b6207536e3.I7d170a48a9096e6b7269c3a9f447c
326f929b171@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b826beaf302033e4d4b483006be7284f47bb2521
Author: Eric Dumazet <edumazet@google.com>
Date:   Tue Jan 18 03:43:40 2022 -0800

    netns: add schedule point in ops_exit_list()

    commit 2836615aa22de55b8fca5e32fe1b27a67cda625e upstream.

    When under stress, cleanup_net() can have to dismantle
    netns in big numbers. ops_exit_list() currently calls
    many helpers [1] that have no schedule point, and we can
    end up with soft lockups, particularly on hosts
    with many cpus.

    Even for moderate amount of netns processed by cleanup_net()
    this patch avoids latency spikes.

    [1] Some of these helpers like fib_sync_up() and fib_sync_down_dev()
    are very slow because net/ipv4/fib_semantics.c uses host-wide hash tables,
    and ifindex is used as the only input of two hash functions.
        ifindexes tend to be the same for all netns (lo.ifindex==1 per instance)
        This will be fixed in a separate patch.

    Fixes: 72ad937abd0a ("net: Add support for batching network namespace cleanups")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Eric W. Biederman <ebiederm@xmission.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8ae336438b3d0b5608f673174edf97c9f1b0ac37
Author: Eric Dumazet <edumazet@google.com>
Date:   Thu Jan 13 01:22:29 2022 -0800

    inet: frags: annotate races around fqdir->dead and fqdir->high_thresh

    commit 91341fa0003befd097e190ec2a4bf63ad957c49a upstream.

    Both fields can be read/written without synchronization,
    add proper accessors and documentation.
```

```
    Fixes: d5dd88794a13 ("inet: fix various use-after-free in defrags units")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7d353ec54d2586d8cc94316dcc88b8760e789351
Author: Eric W. Biederman <ebiederm@xmission.com>
Date:   Mon Jan 3 11:32:36 2022 -0600

    taskstats: Cleanup the use of task->exit_code

    commit 1b5a42d9c85f0e731f01c8d1129001fd8531a8a0 upstream.

    In the function bacct_add_task the code reading task->exit_code was
    introduced in commit f3cef7a99469 ("[PATCH] csa: basic accounting over
    taskstats"), and it is not entirely clear what the taskstats interface
    is trying to return as only returning the exit_code of the first task
    in a process doesn't make a lot of sense.

    As best as I can figure the intent is to return task->exit_code after
    a task exits.  The field is returned with per task fields, so the
    exit_code of the entire process is not wanted.  Only the value of the
    first task is returned so this is not a useful way to get the per task
    ptrace stop code.  The ordinary case of returning this value is
    returning after a task exits, which also precludes use for getting
    a ptrace value.

    It is common to for the first task of a process to also be the last
    task of a process so this field may have done something reasonable by
    accident in testing.

    Make ac_exitcode a reliable per task value by always returning it for
    every exited task.

    Setting ac_exitcode in a sensible mannter makes it possible to continue
    to provide this value going forward.

    Cc: Balbir Singh <bsingharora@gmail.com>
    Fixes: f3cef7a99469 ("[PATCH] csa: basic accounting over taskstats")
    Link: https://lkml.kernel.org/r/20220103213312.9144-5-ebiederm@xmission.com
    Signed-off-by: "Eric W. Biederman" <ebiederm@xmission.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 45a05e32d6eca1bd65ab93a63f324fd556937ea4
Author: Michael S. Tsirkin <mst@redhat.com>
Date:   Thu Jan 6 07:57:46 2022 -0500

    virtio_ring: mark ring unused on error

    commit 1861ba626ae9b98136f3e504208cdef6b29cd3ec upstream.

    A recently added error path does not mark ring unused when exiting on
    OOM, which will lead to BUG on the next entry in debug builds.

    TODO: refactor code so we have START_USE and END_USE in the same function.

    Fixes: fc6d70f40b3d ("virtio_ring: check desc == NULL when using indirect with packed")
    Cc: "Xuan Zhuo" <xuanzhuo@linux.alibaba.com>
    Cc: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Xuan Zhuo <xuanzhuo@linux.alibaba.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bd33ff5842c63ebb831b4fdc18d61c1805329f00
Author: Eli Cohen <elic@nvidia.com>
```

```
Date:   Thu Dec 30 16:20:24 2021 +0200

    vdpa/mlx5: Fix wrong configuration of virtio_version_1_0

    commit 97143b70aa847f2b0a1f959dde126b76ff7b5376 upstream.

    Remove overriding of virtio_version_1_0 which forced the virtqueue
    object to version 1.

    Fixes: 1a86b377aa21 ("vdpa/mlx5: Add VDPA driver for supported mlx5 devices")
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20211230142024.142979-1-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Reviewed-by: Parav Pandit <parav@nvidia.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Reviewed-by: Si-Wei Liu <si-wei.liu@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4ad1232ab35c8d7c892965a88fa24de19c299eae
Author: Laurence de Bruxelles <lfdebrux@gmail.com>
Date:   Sat Jan 1 15:41:49 2022 +0000

    rtc: pxa: fix null pointer dereference

    commit 34127b3632b21e5c391756e724b1198eb9917981 upstream.

    With the latest stable kernel versions the rtc on the PXA based
    Zaurus does not work, when booting I see the following kernel messages:

    pxa-rtc pxa-rtc: failed to find rtc clock source
    pxa-rtc pxa-rtc: Unable to init SA1100 RTC sub-device
    pxa-rtc: probe of pxa-rtc failed with error -2
    hctosys: unable to open rtc device (rtc0)

    I think this is because commit f2997775b111 ("rtc: sa1100: fix possible
    race condition") moved the allocation of the rtc_device struct out of
    sa1100_rtc_init and into sa1100_rtc_probe. This means that pxa_rtc_probe
    also needs to do allocation for the rtc_device struct, otherwise
    sa1100_rtc_init will try to dereference a null pointer. This patch adds
    that allocation by copying how sa1100_rtc_probe in
    drivers/rtc/rtc-sa1100.c does it; after the IRQs are set up a managed
    rtc_device is allocated.

    I've tested this patch with `qemu-system-arm -machine akita` and with a
    real Zaurus SL-C1000 applied to 4.19, 5.4, and 5.10.

    Signed-off-by: Laurence de Bruxelles <lfdebrux@gmail.com>
    Fixes: f2997775b111 ("rtc: sa1100: fix possible race condition")
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/20220101154149.12026-1-lfdebrux@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb27f792e4b73938b6dfb2fd33ebcb7bb59d2717
Author: Kees Cook <keescook@chromium.org>
Date:   Wed Dec 8 20:39:15 2021 -0800

    rtc: Move variable into switch case statement

    commit ba52eac083e1598e748811ff58d259f77e4c5c4d upstream.

    When building with automatic stack variable initialization, GCC 12
    complains about variables defined outside of switch case statements.
    Move the variable into the case that uses it, which silences the warning:

    drivers/rtc/dev.c: In function 'rtc_dev_ioctl':
```

```
    drivers/rtc/dev.c:394:30: warning: statement will never be executed [-Wswitch-
unreachable]
      394 |                         long offset;
          |                              ^~~~~~
```

    Fixes: 6a8af1b6568a ("rtc: add parameter ioctl")
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Reviewed-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/20211209043915.1378393-1-keescook@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1164ad5937880012e109c1f4db3020e64147b8f4
Author: Matt Johnston <matt@codeconstruct.com.au>
Date: Mon Jan 10 10:18:06 2022 +0800

    mctp: test: zero out sockaddr

    commit 284a4d94e8e74fbd731ee67e29196656ca823423 upstream.

    MCTP now requires that padding bytes are zero.

    Signed-off-by: Matt Johnston <matt@codeconstruct.com.au>
    Fixes: 1e4b50f06d97 ("mctp: handle the struct sockaddr_mctp padding fields")
    Link: https://lore.kernel.org/r/20220110021806.2343023-1-matt@codeconstruct.com.au
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ff027c3aec166a5752df36605c09918bbf82ae82
Author: Dmitry Torokhov <dmitry.torokhov@gmail.com>
Date: Fri Jan 7 12:09:36 2022 -0800

    HID: vivaldi: fix handling devices not using numbered reports

    commit 3fe6acd4dc922237b30e55473c9349c6ce0690f3 upstream.

    Unfortunately details of USB HID transport bled into HID core and
    handling of numbered/unnumbered reports is quite a mess, with
    hid_report_len() calculating the length according to USB rules,
    and hid_hw_raw_request() adding report ID to the buffer for both
    numbered and unnumbered reports.

    Untangling it all requres a lot of changes in HID, so for now let's
    handle this in the driver.

    [jkosina@suse.cz: microoptimize field->report->id to report->id]
    Fixes: 14c9c014babe ("HID: add vivaldi HID driver")
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Tested-by: Stephen Boyd <swboyd@chromium.org> # CoachZ
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dd53581ee95b0ff09226360f748d28110e02f1e9
Author: Johannes Berg <johannes.berg@intel.com>
Date: Mon Sep 13 10:17:29 2021 +0200

    um: gitignore: Add kernel/capflags.c

    commit 4b86366fdfbedec42f8f7ee037775f2839921d34 upstream.

    This file is generated, we should ignore it.

    Fixes: d8fb32f4790f ("um: Add support for host CPU flags and alignment")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Acked-By: anton.ivanov@cambridgegreys.com

```
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d08e39865d665917352380862dacc41c485d0efa
Author: Yury Norov <yury.norov@gmail.com>
Date:   Sat Aug 14 14:16:57 2021 -0700

    bitops: protect find_first_{,zero}_bit properly

    commit b7ec62d7ee0f0b8af6ba190501dff7f9ee6545ca upstream.

    find_first_bit() and find_first_zero_bit() are not protected with
    ifdefs as other functions in find.h. It causes build errors on some
    platforms if CONFIG_GENERIC_FIND_FIRST_BIT is enabled.

    Signed-off-by: Yury Norov <yury.norov@gmail.com>
    Fixes: 2cc7b6a44ac2 ("lib: add fast path for find_first_*_bit() and find_last_bit()")
    Reported-by: kernel test robot <lkp@intel.com>
    Tested-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8f28893130729439ca277ed3d541e542f46cfaa2
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:32 2022 -0600

    net: axienet: increase default TX ring size to 128

    commit 2d19c3fd80178160dd505ccd7fed1643831227a5 upstream.

    With previous changes to make the driver handle the TX ring size more
    correctly, the default TX ring size of 64 appears to significantly
    bottleneck TX performance to around 600 Mbps on a 1 Gbps link on ZynqMP.
    Increasing this to 128 seems to bring performance up to near line rate and
    shouldn't cause excess bufferbloat (this driver doesn't yet support modern
    byte-based queue management).

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 24727a75358d01058c0a852be1ca94b7fb771456
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:31 2022 -0600

    net: axienet: fix for TX busy handling

    commit bb193e3db8b86a63f26889c99e14fd30c9ebd72a upstream.

    Network driver documentation indicates we should be avoiding returning
    NETDEV_TX_BUSY from ndo_start_xmit in normal cases, since it requires
    the packets to be requeued. Instead the queue should be stopped after
    a packet is added to the TX ring when there may not be enough room for an
    additional one. Also, when TX ring entries are completed, we should only
    wake the queue if we know there is room for another full maximally
    fragmented packet.

    Print a warning if there is insufficient space at the start of start_xmit,
    since this should no longer happen.

    Combined with increasing the default TX ring size (in a subsequent
    patch), this appears to recover the TX performance lost by previous changes
    to actually manage the TX ring state properly.

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
```

```
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 821872d2830eca2b6636245a10b9059d09382263
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:30 2022 -0600

    net: axienet: fix number of TX ring slots for available check

    commit aba57a823d2985a2cc8c74a2535f3a88e68d9424 upstream.

    The check for the number of available TX ring slots was off by 1 since a
    slot is required for the skb header as well as each fragment. This could
    result in overwriting a TX ring slot that was still in use.

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 69ddcc3faac72f80802843fca54a31f39a9a8609
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:29 2022 -0600

    net: axienet: Fix TX ring slot available check

    commit 996defd7f8b5dafc1d480b7585c7c62437f80c3c upstream.

    The check for whether a TX ring slot was available was incorrect,
    since a slot which had been loaded with transmit data but the device had
    not started transmitting would be treated as available, potentially
    causing non-transmitted slots to be overwritten. The control field in
    the descriptor should be checked, rather than the status field (which may
    only be updated when the device completes the entry).

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b910283f27ff0acc3596e085879292397cd2ad83
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:28 2022 -0600

    net: axienet: limit minimum TX ring size

    commit 70f5817deddbc6ef3faa35841cab83c280cc653a upstream.

    The driver will not work properly if the TX ring size is set to below
    MAX_SKB_FRAGS + 1 since it needs to hold at least one full maximally
    fragmented packet in the TX ring. Limit setting the ring size to below
    this value.

    Fixes: 8b09ca823ffb4 ("net: axienet: Make RX/TX ring sizes configurable")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8eaa258de820b9fff3832740285d935634df1440
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:27 2022 -0600

    net: axienet: add missing memory barriers
```

```
    commit 95978df6fa328df619c15312e65ece469c2be2d2 upstream.

    This driver was missing some required memory barriers:

    Use dma_rmb to ensure we see all updates to the descriptor after we see
    that an entry has been completed.

    Use wmb and rmb to avoid stale descriptor status between the TX path and
    TX complete IRQ path.

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 190ab3e09053a934b3f054273770c5a9bf9d0c78
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:26 2022 -0600

    net: axienet: reset core on initialization prior to MDIO access

    commit 04cc2da39698efd7eb2e30c112538922d26f848e upstream.

    In some cases where the Xilinx Ethernet core was used in 1000Base-X or
    SGMII modes, which use the internal PCS/PMA PHY, and the MGT
    transceiver clock source for the PCS was not running at the time the
    FPGA logic was loaded, the core would come up in a state where the
    PCS could not be found on the MDIO bus. To fix this, the Ethernet core
    (including the PCS) should be reset after enabling the clocks, prior to
    attempting to access the PCS using of_mdio_find_device.

    Fixes: 1a02556086fc (net: axienet: Properly handle PCS/PMA PHY for 1000BaseX mode)
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 137f5b22d564ef6202558851559d814717ab21af
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:25 2022 -0600

    net: axienet: Wait for PhyRstCmplt after core reset

    commit b400c2f4f4c53c86594dd57098970d97d488bfde upstream.

    When resetting the device, wait for the PhyRstCmplt bit to be set
    in the interrupt status register before continuing initialization, to
    ensure that the core is actually ready. When using an external PHY, this
    also ensures we do not start trying to access the PHY while it is still
    in reset. The PHY reset is initiated by the core reset which is
    triggered just above, but remains asserted for 5ms after the core is
    reset according to the documentation.

    The MgtRdy bit could also be waited for, but unfortunately when using
    7-series devices, the bit does not appear to work as documented (it
    seems to behave as some sort of link state indication and not just an
    indication the transceiver is ready) so it can't really be relied on for
    this purpose.

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 113c4fc011b76401b5ce9fa0a019db132eacd725
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Tue Jan 18 15:41:24 2022 -0600

    net: axienet: increase reset timeout

    commit 2e5644b1bab2ccea9cfc7a9520af95b94eb0dbf1 upstream.

    The previous timeout of 1ms was too short to handle some cases where the
    core is reset just after the input clocks were started, which will
    be introduced in an upcoming patch. Increase the timeout to 50ms. Also
    simplify the reset timeout checking to use read_poll_timeout.

    Fixes: 8a3b7a252dca9 ("drivers/net/ethernet/xilinx: added Xilinx AXI Ethernet driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 78148e0b14b48ed97e65ea2eebadb394f3e41f1b
Author: Wen Gu <guwen@linux.alibaba.com>
Date:   Sun Jan 16 15:43:42 2022 +0800

    net/smc: Fix hung_task when removing SMC-R devices

    commit 56d99e81ecbc997a5f984684d0eeb583992b2072 upstream.

    A hung_task is observed when removing SMC-R devices. Suppose that
    a link group has two active links(lnk_A, lnk_B) associated with two
    different SMC-R devices(dev_A, dev_B). When dev_A is removed, the
    link group will be removed from smc_lgr_list and added into
    lgr_linkdown_list. lnk_A will be cleared and smcibdev(A)->lnk_cnt
    will reach to zero. However, when dev_B is removed then, the link
    group can't be found in smc_lgr_list and lnk_B won't be cleared,
    making smcibdev->lnk_cnt never reaches zero, which causes a hung_task.

    This patch fixes this issue by restoring the implementation of
    smc_smcr_terminate_all() to what it was before commit 349d43127dac
    ("net/smc: fix kernel panic caused by race of smc_sock"). The original
    implementation also satisfies the intention that make sure QP destroy
    earlier than CQ destroy because we will always wait for smcibdev->lnk_cnt
    reaches zero, which guarantees QP has been destroyed.

    Fixes: 349d43127dac ("net/smc: fix kernel panic caused by race of smc_sock")
    Signed-off-by: Wen Gu <guwen@linux.alibaba.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5a1687caf19221e8f92737f32d14ed4b32e21723
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Jan 14 06:51:24 2022 +0000

    gpio: idt3243x: Fix IRQ check in idt_gpio_probe

    commit 30fee1d7462a446ade399c0819717a830cbdca69 upstream.

    platform_get_irq() returns negative error number instead 0 on failure.
    And the doc of platform_get_irq() provides a usage example:

        int irq = platform_get_irq(pdev, 0);
        if (irq < 0)
            return irq;

    Fix the check of return value to catch errors correctly.
```

```
    Fixes: 4195926aedca ("gpio: Add support for IDT 79RC3243x GPIO controller")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b8649c9b78f9231b1e290a18288b7160b2152b96
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Jan 14 06:48:20 2022 +0000

    gpio: mpc8xxx: Fix IRQ check in mpc8xxx_probe

    commit 0b39536cc699db6850c426db7f9cb45923de40c5 upstream.

    platform_get_irq() returns negative error number instead 0 on failure.
    And the doc of platform_get_irq() provides a usage example:

        int irq = platform_get_irq(pdev, 0);
        if (irq < 0)
            return irq;

    Fix the check of return value to catch errors correctly.

    Fixes: 76c47d1449fc ("gpio: mpc8xxx: Add ACPI support")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 685c416ba439c1bf130da7626272555614a8cf65
Author: John Keeping <john@metanate.com>
Date:   Fri Nov 26 15:13:52 2021 +0000

    pinctrl/rockchip: fix gpio device creation

    commit bceb6732f3fd2a55d8f2e518cced1c7555e216b6 upstream.

    GPIO nodes are not themselves busses, so passing rockchip_bank_match
    here is wrong.  Passing NULL instead uses the standard bus match table
    which is more appropriate.

    devm_of_platform_populate() shows that this is the normal way to call
    of_platform_populate() from a device driver, so in order to match that
    more closely also add the pinctrl device as the parent for the newly
    created GPIO controllers.

    Specifically, using the wrong match here can break dynamic GPIO hogs as
    marking the GPIO bank as a bus means that of_platform_notify() will set
    OF_POPULATED on new child nodes and if this happens before
    of_gpio_notify() is called then the new hog will be skipped as
    OF_POPULATED is already set.

    Fixes: 9ce9a02039de ("pinctrl/rockchip: drop the gpio related codes")
    Signed-off-by: John Keeping <john@metanate.com>
    Link: https://lore.kernel.org/r/20211126151352.1509583-1-john@metanate.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a1c0e376b00416721ff2e199cdc2a980730e08c3
Author: Robert Hancock <robert.hancock@calian.com>
Date:   Wed Jan 12 14:38:16 2022 -0600

    clk: si5341: Fix clock HW provider cleanup

    commit 49a8f2bc8d88702783c7e163ec84374e9a022f71 upstream.
```

```
    The call to of_clk_add_hw_provider was not undone on remove or on probe
    failure, which could cause an oops on a subsequent attempt to retrieve
    clocks for the removed device. Switch to the devm version of the
    function to avoid this issue.

    Fixes: 3044a860fd09 ("clk: Add Si5341/Si5340 driver")
    Signed-off-by: Robert Hancock <robert.hancock@calian.com>
    Link: https://lore.kernel.org/r/20220112203816.1784610-1-robert.hancock@calian.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 68587f89494b443b3461e2af50811e868f6759e9
Author: Stephen Boyd <sboyd@kernel.org>
Date:   Thu Dec 9 17:34:05 2021 -0800

    clk: Emit a stern warning with writable debugfs enabled

    commit 489a71964f9d74e697a12cd0ace20ed829eb1f93 upstream.

    We don't want vendors to be enabling this part of the clk code and
    shipping it to customers. Exposing the ability to change clk frequencies
    and parents via debugfs is potentially damaging to the system if folks
    don't know what they're doing. Emit a strong warning so that the message
    is clear: don't enable this outside of development systems.

    Fixes: 37215da5553e ("clk: Add support for setting clk_rate via debugfs")
    Cc: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20211210014237.2130300-1-sboyd@kernel.org
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 47c99508c325688e5851afe58b3bcb39be091211
Author: Eric Dumazet <edumazet@google.com>
Date:   Fri Jan 14 08:43:28 2022 -0800

    af_unix: annote lockless accesses to unix_tot_inflight & gc_in_progress

    commit 9d6d7f1cb67cdee15f1a0e85aacfb924e0e02435 upstream.

    wait_for_unix_gc() reads unix_tot_inflight & gc_in_progress
    without synchronization.

    Adds READ_ONCE()/WRITE_ONCE() and their associated comments
    to better document the intent.

    BUG: KCSAN: data-race in unix_inflight / wait_for_unix_gc

    write to 0xffffffff86e2b7c0 of 4 bytes by task 9380 on cpu 0:
     unix_inflight+0x1e8/0x260 net/unix/scm.c:63
     unix_attach_fds+0x10c/0x1e0 net/unix/scm.c:121
     unix_scm_to_skb net/unix/af_unix.c:1674 [inline]
     unix_dgram_sendmsg+0x679/0x16b0 net/unix/af_unix.c:1817
     unix_seqpacket_sendmsg+0xcc/0x110 net/unix/af_unix.c:2258
     sock_sendmsg_nosec net/socket.c:704 [inline]
     sock_sendmsg net/socket.c:724 [inline]
     ____sys_sendmsg+0x39a/0x510 net/socket.c:2409
     ___sys_sendmsg net/socket.c:2463 [inline]
     __sys_sendmmsg+0x267/0x4c0 net/socket.c:2549
     __do_sys_sendmmsg net/socket.c:2578 [inline]
     __se_sys_sendmmsg net/socket.c:2575 [inline]
     __x64_sys_sendmmsg+0x53/0x60 net/socket.c:2575
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x44/0xd0 arch/x86/entry/common.c:80
     entry_SYSCALL_64_after_hwframe+0x44/0xae
```

```
    read to 0xffffffff86e2b7c0 of 4 bytes by task 9375 on cpu 1:
     wait_for_unix_gc+0x24/0x160 net/unix/garbage.c:196
     unix_dgram_sendmsg+0x8e/0x16b0 net/unix/af_unix.c:1772
     unix_seqpacket_sendmsg+0xcc/0x110 net/unix/af_unix.c:2258
     sock_sendmsg_nosec net/socket.c:704 [inline]
     sock_sendmsg net/socket.c:724 [inline]
     ____sys_sendmsg+0x39a/0x510 net/socket.c:2409
     ___sys_sendmsg net/socket.c:2463 [inline]
     __sys_sendmmsg+0x267/0x4c0 net/socket.c:2549
     __do_sys_sendmmsg net/socket.c:2578 [inline]
     __se_sys_sendmmsg net/socket.c:2575 [inline]
     __x64_sys_sendmmsg+0x53/0x60 net/socket.c:2575
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x44/0xd0 arch/x86/entry/common.c:80
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    value changed: 0x00000002 -> 0x00000004

    Reported by Kernel Concurrency Sanitizer on:
    CPU: 1 PID: 9375 Comm: syz-executor.1 Not tainted 5.16.0-rc7-syzkaller #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
01/01/2011

    Fixes: 9915672d4127 ("af_unix: limit unix_tot_inflight")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Link: https://lore.kernel.org/r/20220114164328.2038499-1-eric.dumazet@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 76a01740ada1ed9149e4a6a9d2f3aa4409ca0d66
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Sat Nov 27 17:10:27 2021 +0300

    crypto: octeontx2 - uninitialized variable in kvf_limits_store()

    commit 0ea275df84c389e910a3575a9233075118c173ee upstream.

    If kstrtoint() fails then "lfs_num" is uninitialized and the warning
    doesn't make any sense.  Just delete it.

    Fixes: 8ec8015a3168 ("crypto: octeontx2 - add support to process the crypto request")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit efc0f3f844503d2fd478272039d5f8539a533c09
Author: Chao Yu <chao@kernel.org>
Date:   Sun Dec 12 20:28:12 2021 +0800

    f2fs: fix to check available space of CP area correctly in update_ckpt_flags()

    commit b702c83e2eaa2fa2d72e957c55c0321535cc8b9f upstream.

    Otherwise, nat_bit area may be persisted across boundary of CP area during
    nat_bit rebuilding.

    Fixes: 94c821fb286b ("f2fs: rebuild nat_bits during umount")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fa956a1ddc5398d56fa03e00ba1c8cdd3ef9b827
Author: Chao Yu <chao@kernel.org>
```

```
Date:   Sat Dec 11 21:27:36 2021 +0800

    f2fs: fix to reserve space for IO align feature

    commit 300a842937fbcfb5a189cea9ba15374fdb0b5c6b upstream.

    https://bugzilla.kernel.org/show_bug.cgi?id=204137

    With below script, we will hit panic during new segment allocation:

    DISK=bingo.img
    MOUNT_DIR=/mnt/f2fs

    dd if=/dev/zero of=$DISK bs=1M count=105
    mkfs.f2fe -a 1 -o 19 -t 1 -z 1 -f -q $DISK

    mount -t f2fs $DISK $MOUNT_DIR -o
"noinline_dentry,flush_merge,noextent_cache,mode=lfs,io_bits=7,fsync_mode=strict"

    for (( i = 0; i < 4096; i++ )); do
            name=`head /dev/urandom | tr -dc A-Za-z0-9 | head -c 10`
            mkdir $MOUNT_DIR/$name
    done

    umount $MOUNT_DIR
    rm $DISK

commit ad8ea3747d8a805def0fd71b1c56eaa38068de31
Author: Hyeong-Jun Kim <hj514.kim@samsung.com>
Date:   Fri Dec 10 13:30:12 2021 +0900

    f2fs: compress: fix potential deadlock of compress file

    commit 7377e853967ba45bf409e3b5536624d2cbc99f21 upstream.

    There is a potential deadlock between writeback process and a process
    performing write_begin() or write_cache_pages() while trying to write
    same compress file, but not compressable, as below:

    [Process A] - doing checkpoint
    [Process B]                     [Process C]
    f2fs_write_cache_pages()
    - lock_page() [all pages in cluster, 0-31]
    - f2fs_write_multi_pages()
     - f2fs_write_raw_pages()
      - f2fs_write_single_data_page()
       - f2fs_do_write_data_page()
         - return -EAGAIN [f2fs_trylock_op() failed]
       - unlock_page(page) [e.g., page 0]
                                    - generic_perform_write()
                                     - f2fs_write_begin()
                                       - f2fs_prepare_compress_overwrite()
                                        - prepare_compress_overwrite()
                                          - lock_page() [e.g., page 0]
                                          - lock_page() [e.g., page 1]
       - lock_page(page) [e.g., page 0]

    Since there is no compress process, it is no longer necessary to hold
    locks on every pages in cluster within f2fs_write_raw_pages().

    This patch changes f2fs_write_raw_pages() to release all locks first
    and then perform write same as the non-compress file in
    f2fs_write_cache_pages().

    Fixes: 4c8ff7095bef ("f2fs: support data compression")
```

```
    Signed-off-by: Hyeong-Jun Kim <hj514.kim@samsung.com>
    Signed-off-by: Sungjong Seo <sj1557.seo@samsung.com>
    Signed-off-by: Youngjin Gil <youngjin.gil@samsung.com>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c50443a829054989fc165cd0003a1493e16f185e
Author: Chao Yu <chao@kernel.org>
Date:   Mon Dec 6 22:44:20 2021 +0800

    f2fs: fix to avoid panic in is_alive() if metadata is inconsistent

    commit f6db43076d190d9bf75559dec28e18b9d12e4ce5 upstream.

    As report by Wenqing Liu in bugzilla:

    https://bugzilla.kernel.org/show_bug.cgi?id=215231

    If we enable CONFIG_F2FS_CHECK_FS config, and with fuzzed image attached
    in above link, we will encounter panic when executing below script:

    1. mkdir mnt
    2. mount -t f2fs tmp1.img mnt
    3. touch tmp

    F2FS-fs (loop11): mismatched blkaddr 5765 (source_blkaddr 1) in seg 3
    kernel BUG at fs/f2fs/gc.c:1042!
     do_garbage_collect+0x90f/0xa80 [f2fs]
     f2fs_gc+0x294/0x12a0 [f2fs]
     f2fs_balance_fs+0x2c5/0x7d0 [f2fs]
     f2fs_create+0x239/0xd90 [f2fs]
     lookup_open+0x45e/0xa90
     open_last_lookups+0x203/0x670
     path_openat+0xae/0x490
     do_filp_open+0xbc/0x160
     do_sys_openat2+0x2f1/0x500
     do_sys_open+0x5e/0xa0
     __x64_sys_openat+0x28/0x40

    Previously, f2fs tries to catch data inconcistency exception in between
    SSA and SIT table during GC, however once the exception is caught, it will
    call f2fs_bug_on to hang kernel, it's not needed, instead, let's set
    SBI_NEED_FSCK flag and skip migrating current block.

    Fixes: bbf9f7d90f21 ("f2fs: Fix indefinite loop in f2fs_gc()")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cbb6ce0ffc1ce8cfd603fd83bf599807c61c0d94
Author: Fengnan Chang <changfengnan@vivo.com>
Date:   Fri Nov 26 18:19:19 2021 +0800

    f2fs: fix remove page failed in invalidate compress pages

    commit d1917865a7906baf6b687e15e8e6195a295a3992 upstream.

    Since compress inode not a regular file, generic_error_remove_page in
    f2fs_invalidate_compress_pages will always be failed, set compress
    inode as a regular file to fix it.

    Fixes: 6ce19aff0b8c ("f2fs: compress: add compress_inode to cache compressed blocks")
    Signed-off-by: Fengnan Chang <changfengnan@vivo.com>
    Reviewed-by: Chao Yu <chao@kernel.org>
```

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

commit 38624fecc7f782871b47738a05df4d2dc6eb3e98
Author: Zack Rusin <zackr@vmware.com>
Date:   Wed Dec 15 13:41:47 2021 -0500

    drm/vmwgfx: Remove unused compile options

    commit 50ca8cc7c0fdd9ab16b8b66ffb301fface101fac upstream.

    Before the driver had screen targets support we had to disable explicit
    bringup of its infrastructure because it was breaking screen objects
    support.
    Since the implementation of screen targets landed there hasn't been a
    reason to explicitly disable it and the options were never used.
    Remove of all that unused code.

    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Fixes: d80efd5cb3de ("drm/vmwgfx: Initial DX support")
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211215184147.3688785-3-
zack@kde.org
    (cherry picked from commit 11343099d5ae6c7411da1425b6b162c89fb5bf10)
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit faa77dc9fc4b50185e8ca9b3f7f85f9c237f3f48
Author: Zack Rusin <zackr@vmware.com>
Date:   Wed Dec 15 13:41:46 2021 -0500

    drm/vmwgfx: Remove explicit transparent hugepages support

    commit bc701a28c74e78d7b5aa2b8628cb3608d4785d14 upstream.

    Old versions of the svga device used to export virtual vram, handling of
    which was optimized on top of transparent hugepages support. Only very
    old devices (OpenGL 2.1 support and earlier) used this code and at this
    point performance differences are negligible.

    Because the code requires very old hardware versions to run it has
    been largely untested and unused for a long time.

    Furthermore removal of the ttm hugepages support in:
    commit 0d979509539e ("drm/ttm: remove ttm_bo_vm_insert_huge()")
    broke the coherency mode in vmwgfx when running with hugepages.

    Fixes: 0d979509539e ("drm/ttm: remove ttm_bo_vm_insert_huge()")
    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Cc: Jason Gunthorpe <jgg@nvidia.com>
    Cc: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: Daniel Vetter <daniel.vetter@ffwll.ch>
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Reviewed-by: Maaz Mombasawala <mombasawalam@vmware.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211215184147.3688785-2-
zack@kde.org
    (cherry picked from commit 49d535d64d52945e2c874f380705675e20a02b6a)
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4adf1e946ab54d007c74a0be3e75d9b37ce16ba5
Author: Geert Uytterhoeven <geert@linux-m68k.org>
Date:   Fri Dec 17 13:49:24 2021 +0100
```

```
      riscv: dts: microchip: mpfs: Drop empty chosen node

      commit 53ef07326ad0d6ae7fefded22bc53b427d542761 upstream.

      It does not make sense to have an (empty) chosen node in an SoC-specific
      .dtsi, as chosen is meant for system-specific configuration.
      It is already provided in microchip-mpfs-icicle-kit.dts anyway.

      Fixes: 0fa6107eca4186ad ("RISC-V: Initial DTS for Microchip ICICLE board")
      Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
      Reviewed-by: Conor Dooley <conor.dooley@microchip.com>
      Tested-by: Conor Dooley <conor.dooley@microchip.com>
      Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
      Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit ef0212bc1c6d087480ee069890c99e40f5131071
  Author: Palmer Dabbelt <palmer@rivosinc.com>
  Date:   Fri Nov 19 08:44:02 2021 -0800

      RISC-V: defconfigs: Set CONFIG_FB=y, for FB console

      commit 3d12b634fe8206ea974c6061a3f3eea529ffbc48 upstream.

      We have CONFIG_FRAMEBUFFER_CONSOLE=y in the defconfigs, but that depends
      on CONFIG_FB so it's not actually getting set.  I'm assuming most users
      on real systems want a framebuffer console, so this enables CONFIG_FB to
      allow that to take effect.

      Fixes: 33c57c0d3c67 ("RISC-V: Add a basic defconfig")
      Reviewed-by: Anup Patel <anup@brainfault.org>
      Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit dcb684740f0b8fd7fc6df84a00b7d6f1bc50f01a
  Author: Miaoqian Lin <linmq006@gmail.com>
  Date:   Thu Jan 20 12:18:12 2022 +0000

      parisc: pdc_stable: Fix memory leak in pdcs_register_pathentries

      commit d24846a4246b6e61ecbd036880a4adf61681d241 upstream.

      kobject_init_and_add() takes reference even when it fails.
      According to the doc of kobject_init_and_add() :

         If this function returns an error, kobject_put() must be called to
         properly clean up the memory associated with the object.

      Fix memory leak by calling kobject_put().

      Fixes: 73f368cf679b ("Kobject: change drivers/parisc/pdc_stable.c to use
  kobject_init_and_add")
      Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
      Signed-off-by: Helge Deller <deller@gmx.de>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 852392448b5db0e4470bcfe7a76a70834651696b
  Author: Tobias Waldekranz <tobias@waldekranz.com>
  Date:   Tue Jan 18 22:50:53 2022 +0100

      net/fsl: xgmac_mdio: Fix incorrect iounmap when removing module

      commit 3f7c239c7844d2044ed399399d97a5f1c6008e1b upstream.

      As reported by sparse: In the remove path, the driver would attempt to
```

```
    unmap its own priv pointer - instead of the io memory that it mapped
    in probe.

    Fixes: 9f35a7342cff ("net/fsl: introduce Freescale 10G MDIO driver")
    Signed-off-by: Tobias Waldekranz <tobias@waldekranz.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ebe929e5d3fdd8c965fa562dcf460b8f2fc52768
Author: Tobias Waldekranz <tobias@waldekranz.com>
Date:   Tue Jan 18 22:50:50 2022 +0100

    net/fsl: xgmac_mdio: Add workaround for erratum A-009885

    commit 6198c722019774d38018457a8bfb9ba3ed8c931e upstream.

    Once an MDIO read transaction is initiated, we must read back the data
    register within 16 MDC cycles after the transaction completes. Outside
    of this window, reads may return corrupt data.

    Therefore, disable local interrupts in the critical section, to
    maximize the probability that we can satisfy this requirement.

    Fixes: d55ad2967d89 ("powerpc/mpc85xx: Create dts components for the FSL QorIQ DPAA
FMan")
    Signed-off-by: Tobias Waldekranz <tobias@waldekranz.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 42f62354edb76f80b7cf4f234b60ecf70a7dde9e
Author: Guillaume Nault <gnault@redhat.com>
Date:   Mon Jan 10 14:43:14 2022 +0100

    mlx5: Don't accidentally set RTO_ONLINK before mlx5e_route_lookup_ipv4_get()

    commit 48d67543e01d73292e0bb66d3f10fc422e79e031 upstream.

    Mask the ECN bits before calling mlx5e_route_lookup_ipv4_get(). The
    tunnel key might have the last ECN bit set. This interferes with the
    route lookup process as ip_route_output_key_hash() interpretes this bit
    specially (to restrict the route scope).

    Found by code inspection, compile tested only.

    Fixes: c7b9038d8af6 ("net/mlx5e: TC preparation refactoring for routing update event")
    Fixes: 9a941117fb76 ("net/mlx5e: Maximize ip tunnel key usage on the TC offloading
path")
    Signed-off-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b6116c0fedb2abe4eee3ff15a6a9a20ba8e6e950
Author: Eric Dumazet <edumazet@google.com>
Date:   Wed Jan 19 02:04:12 2022 -0800

    ipv4: avoid quadratic behavior in netns dismantle

    commit d07418afea8f1d9896aaf9dc5ae47ac4f45b220c upstream.

    net/ipv4/fib_semantics.c uses an hash table of 256 slots,
    keyed by device ifindexes: fib_info_devhash[DEVINDEX_HASHSIZE]

    Problem is that with network namespaces, devices tend
```

```
    to use the same ifindex.

    lo device for instance has a fixed ifindex of one,
    for all network namespaces.

    This means that hosts with thousands of netns spend
    a lot of time looking at some hash buckets with thousands
    of elements, notably at netns dismantle.

    Simply add a per netns perturbation (net_hash_mix())
    to spread elements more uniformely.

    Also change fib_devindex_hashfn() to use more entropy.

    Fixes: aa79e66eee5d ("net: Make ifindex generation per-net namespace")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reviewed-by: David Ahern <dsahern@kernel.org>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d0334b71cafa953918881e3f1370a515c6e172b3
Author: Eric Dumazet <edumazet@google.com>
Date:   Sun Jan 16 01:02:20 2022 -0800

    ipv4: update fib_info_cnt under spinlock protection

    commit 0a6e6b3c7db6c34e3d149f09cd714972f8753e3f upstream.

    In the past, free_fib_info() was supposed to be called
    under RTNL protection.

    This eventually was no longer the case.

    Instead of enforcing RTNL it seems we simply can
    move fib_info_cnt changes to occur when fib_info_lock
    is held.

    v2: David Laight suggested to update fib_info_cnt
    only when an entry is added/deleted to/from the hash table,
    as fib_info_cnt is used to make sure hash table size
    is optimal.

    BUG: KCSAN: data-race in fib_create_info / free_fib_info

    write to 0xffffffff86e243a0 of 4 bytes by task 26429 on cpu 0:
     fib_create_info+0xe78/0x3440 net/ipv4/fib_semantics.c:1428
     fib_table_insert+0x148/0x10c0 net/ipv4/fib_trie.c:1224
     fib_magic+0x195/0x1e0 net/ipv4/fib_frontend.c:1087
     fib_add_ifaddr+0xd0/0x2e0 net/ipv4/fib_frontend.c:1109
     fib_netdev_event+0x178/0x510 net/ipv4/fib_frontend.c:1466
     notifier_call_chain kernel/notifier.c:83 [inline]
     raw_notifier_call_chain+0x53/0xb0 kernel/notifier.c:391
     __dev_notify_flags+0x1d3/0x3b0
     dev_change_flags+0xa2/0xc0 net/core/dev.c:8872
     do_setlink+0x810/0x2410 net/core/rtnetlink.c:2719
     rtnl_group_changelink net/core/rtnetlink.c:3242 [inline]
     __rtnl_newlink net/core/rtnetlink.c:3396 [inline]
     rtnl_newlink+0xb10/0x13b0 net/core/rtnetlink.c:3506
     rtnetlink_rcv_msg+0x745/0x7e0 net/core/rtnetlink.c:5571
     netlink_rcv_skb+0x14e/0x250 net/netlink/af_netlink.c:2496
     rtnetlink_rcv+0x18/0x20 net/core/rtnetlink.c:5589
     netlink_unicast_kernel net/netlink/af_netlink.c:1319 [inline]
     netlink_unicast+0x5fc/0x6c0 net/netlink/af_netlink.c:1345
     netlink_sendmsg+0x726/0x840 net/netlink/af_netlink.c:1921
     sock_sendmsg_nosec net/socket.c:704 [inline]
```

```
       sock_sendmsg net/socket.c:724 [inline]
       ____sys_sendmsg+0x39a/0x510 net/socket.c:2409
       ___sys_sendmsg net/socket.c:2463 [inline]
       __sys_sendmsg+0x195/0x230 net/socket.c:2492
       __do_sys_sendmsg net/socket.c:2501 [inline]
       __se_sys_sendmsg net/socket.c:2499 [inline]
       __x64_sys_sendmsg+0x42/0x50 net/socket.c:2499
       do_syscall_x64 arch/x86/entry/common.c:50 [inline]
       do_syscall_64+0x44/0xd0 arch/x86/entry/common.c:80
       entry_SYSCALL_64_after_hwframe+0x44/0xae

      read to 0xffffffff86e243a0 of 4 bytes by task 31505 on cpu 1:
       free_fib_info+0x35/0x80 net/ipv4/fib_semantics.c:252
       fib_info_put include/net/ip_fib.h:575 [inline]
       nsim_fib4_rt_destroy drivers/net/netdevsim/fib.c:294 [inline]
       nsim_fib4_rt_replace drivers/net/netdevsim/fib.c:403 [inline]
       nsim_fib4_rt_insert drivers/net/netdevsim/fib.c:431 [inline]
       nsim_fib4_event drivers/net/netdevsim/fib.c:461 [inline]
       nsim_fib_event drivers/net/netdevsim/fib.c:881 [inline]
       nsim_fib_event_work+0x15ca/0x2cf0 drivers/net/netdevsim/fib.c:1477
       process_one_work+0x3fc/0x980 kernel/workqueue.c:2298
       process_scheduled_works kernel/workqueue.c:2361 [inline]
       worker_thread+0x7df/0xa70 kernel/workqueue.c:2447
       kthread+0x2c7/0x2e0 kernel/kthread.c:327
       ret_from_fork+0x1f/0x30

      value changed: 0x00000d2d -> 0x00000d2e

      Reported by Kernel Concurrency Sanitizer on:
      CPU: 1 PID: 31505 Comm: kworker/1:21 Not tainted 5.16.0-rc6-syzkaller #0
      Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
  01/01/2011
      Workqueue: events nsim_fib_event_work

      Fixes: 48bb9eb47b27 ("netdevsim: fib: Add dummy implementation for FIB offload")
      Signed-off-by: Eric Dumazet <edumazet@google.com>
      Reported-by: syzbot <syzkaller@googlegroups.com>
      Cc: David Laight <David.Laight@ACULAB.COM>
      Cc: Ido Schimmel <idosch@mellanox.com>
      Cc: Jiri Pirko <jiri@mellanox.com>
      Reviewed-by: Ido Schimmel <idosch@nvidia.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit fa28881a10d455b30f47bf7b8db3635877b54b57
  Author: German Gomez <german.gomez@arm.com>
  Date:   Tue Jan 18 14:40:54 2022 +0000

      perf evsel: Override attr->sample_period for non-libpfm4 events

      commit 3606c0e1a1050d397ad759a62607e419fd8b0ccb upstream.

      A previous patch preventing "attr->sample_period" values from being
      overridden in pfm events changed a related behaviour in arm-spe.

      Before said patch:

        perf record -c 10000 -e arm_spe_0// -- sleep 1

      Would yield an SPE event with period=10000. After the patch, the period
      in "-c 10000" was being ignored because the arm-spe code initializes
      sample_period to a non-zero value.

      This patch restores the previous behaviour for non-libpfm4 events.
```

```
    Fixes: ae5dcc8abe31 (“perf record: Prevent override of attr->sample_period for libpfm4
events”)
    Reported-by: Chase Conklin <chase.conklin@arm.com>
    Signed-off-by: German Gomez <german.gomez@arm.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Ian Rogers <irogers@google.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: John Fastabend <john.fastabend@gmail.com>
    Cc: KP Singh <kpsingh@kernel.org>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Cc: Martin KaFai Lau <kafai@fb.com>
    Cc: Namhyung Kim <namhyung@kernel.org>
    Cc: Song Liu <songliubraving@fb.com>
    Cc: Stephane Eranian <eranian@google.com>
    Cc: Yonghong Song <yhs@fb.com>
    Cc: bpf@vger.kernel.org
    Cc: netdev@vger.kernel.org
    Link: http://lore.kernel.org/lkml/20220118144054.2541-1-german.gomez@arm.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 71aea35be6ac81c895ec441367579678fb68825e
Author: Daniel Borkmann <daniel@iogearbox.net>
Date:   Fri Jan 14 13:58:36 2022 +0000

    bpf: Mark PTR_TO_FUNC register initially with zero offset

    commit d400a6cf1c8a57cdf10f35220ead3284320d85ff upstream.

    Similar as with other pointer types where we use ldimm64, clear the register
    content to zero first, and then populate the PTR_TO_FUNC type and subprogno
    number. Currently this is not done, and leads to reuse of stale register
    tracking data.

    Given for special ldimm64 cases we always clear the register offset, make it
    common for all cases, so it won't be forgotten in future.

    Fixes: 69c087ba6225 ("bpf: Add bpf_for_each_map_elem() helper")
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: John Fastabend <john.fastabend@gmail.com>
    Acked-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 86320b6c11316236890bd95b1816e0a39046ab41
Author: Yafang Shao <laoar.shao@gmail.com>
Date:   Sat Jan 8 13:46:23 2022 +0000

    bpf: Fix mount source show for bpffs

    commit 1e9d74660d4df625b0889e77018f9e94727ceacd upstream.

    We noticed our tc ebpf tools can't start after we upgrade our in-house kernel
    version from 4.19 to 5.10. That is because of the behaviour change in bpffs
    caused by commit d2935de7e4fd ("vfs: Convert bpf to use the new mount API").

    In our tc ebpf tools, we do strict environment check. If the environment is
    not matched, we won't allow to start the ebpf progs. One of the check is whether
    bpffs is properly mounted. The mount information of bpffs in kernel-4.19 and
    kernel-5.10 are as follows:

    - kernel 4.19
    $ mount -t bpf bpffs /sys/fs/bpf
    $ mount -t bpf
    bpffs on /sys/fs/bpf type bpf (rw,relatime)
```

```
    - kernel 5.10
    $ mount -t bpf bpffs /sys/fs/bpf
    $ mount -t bpf
    none on /sys/fs/bpf type bpf (rw,relatime)

    The device name in kernel-5.10 is displayed as none instead of bpffs, then our
    environment check fails. Currently we modify the tools to adopt to the kernel
    behaviour change, but I think we'd better change the kernel code to keep the
    behavior consistent.

    After this change, the mount information will be displayed the same with the
    behavior in kernel-4.19, for example:

    $ mount -t bpf bpffs /sys/fs/bpf
    $ mount -t bpf
    bpffs on /sys/fs/bpf type bpf (rw,relatime)

    Fixes: d2935de7e4fd ("vfs: Convert bpf to use the new mount API")
    Suggested-by: Daniel Borkmann <daniel@iogearbox.net>
    Signed-off-by: Yafang Shao <laoar.shao@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Christian Brauner <christian.brauner@ubuntu.com>
    Cc: David Howells <dhowells@redhat.com>
    Cc: Al Viro <viro@zeniv.linux.org.uk>
    Link: https://lore.kernel.org/bpf/20220108134623.32467-1-laoar.shao@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a1a658ddbedda26712b6a217412ef84183e22977
Author: Toke Høiland-Jørgensen <toke@redhat.com>
Date:   Fri Jan 7 23:11:13 2022 +0100

    xdp: check prog type before updating BPF link

    commit 382778edc8262b7535f00523e9eb22edba1b9816 upstream.

    The bpf_xdp_link_update() function didn't check the program type before
    updating the program, which made it possible to install any program type as
    an XDP program, which is obviously not good. Syzbot managed to trigger this
    by swapping in an LWT program on the XDP hook which would crash in a helper
    call.

    Fix this by adding a check and bailing out if the types don't match.

    Fixes: 026a4c28e1db ("bpf, xdp: Implement LINK_UPDATE for BPF XDP link")
    Reported-by: syzbot+983941aa85af6ded1fd9@syzkaller.appspotmail.com
    Acked-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Toke Høiland-Jørgensen <toke@redhat.com>
    Link: https://lore.kernel.org/r/20220107221115.326171-1-toke@redhat.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b920399406260a587cea9698a02618cffdac606e
Author: Quentin Monnet <quentin@isovalent.com>
Date:   Wed Nov 10 11:46:30 2021 +0000

    bpftool: Fix indent in option lists in the documentation

    commit 986dec18bbf41f50edc2e0aa4ac5ef8e0f64f328 upstream.

    Mixed indentation levels in the lists of options in bpftool's
    documentation produces some unexpected results. For the "bpftool" man
    page, it prints a warning:

        $ make -C bpftool.8
          GEN     bpftool.8
```

```
         <stdin>:26: (ERROR/3) Unexpected indentation.

    For other pages, there is no warning, but it results in a line break
    appearing in the option lists in the generated man pages.

    RST paragraphs should have a uniform indentation level. Let's fix it.

    Fixes: c07ba629df97 ("tools: bpftool: Update and synchronise option list in doc and help
msg")
    Fixes: 8cc8c6357c8f ("tools: bpftool: Document and add bash completion for -L, -B
options")
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20211110114632.24537-5-quentin@isovalent.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b7dfd3f21309f251203b9665cb704e846ac99a1f
Author: Quentin Monnet <quentin@isovalent.com>
Date:   Wed Nov 10 11:46:28 2021 +0000

    bpftool: Remove inclusion of utilities.mak from Makefiles

    commit 48f5aef4c458c19ab337eed8c95a6486cc014aa3 upstream.

    Bpftool's Makefile, and the Makefile for its documentation, both include
    scripts/utilities.mak, but they use none of the items defined in this
    file. Remove the includes.

    Fixes: 71bb428fe2c1 ("tools: bpf: add bpftool")
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20211110114632.24537-3-quentin@isovalent.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit eb10c8af4cd34efdc46ee2ff937cf586842d980c
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Nov 3 15:08:40 2021 -0700

    libbpf: Remove deprecation attribute from struct bpf_prog_prep_result

    commit 5c5edcdebfcf3a95257b0d8ef27a60af0e0ea03a upstream.

    This deprecation annotation has no effect because for struct deprecation
    attribute has to be declared after struct definition. But instead of
    moving it to the end of struct definition, remove it. When deprecation
    will go in effect at libbpf v0.7, this deprecation attribute will cause
    libbpf's own source code compilation to trigger deprecation warnings,
    which is unavoidable because libbpf still has to support that API.

    So keep deprecation of APIs, but don't mark structs used in API as
    deprecated.

    Fixes: e21d585cb3db ("libbpf: Deprecate multi-instance bpf_program APIs")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Dave Marchevsky <davemarchevsky@fb.com>
    Link: https://lore.kernel.org/bpf/20211103220845.2676888-8-andrii@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 066d48c7b94a0030120b6f4ee84372b1d4c57533
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Oct 25 16:11:07 2021 +0200

    drm/vc4: crtc: Copy assigned channel to the CRTC
```

```
    commit eeb6ab4639590130d25670204ab7b6011333d685 upstream.

    Accessing the crtc->state pointer from outside the modesetting context
    is not allowed. We thus need to copy whatever we need from the KMS state
    to our structure in order to access it.

    In VC4, a number of users of that pointers have crept in over the years,
    and the previous commits removed them all but the HVS channel a CRTC has
    been assigned.

    Let's move this channel in struct vc4_crtc at atomic_begin() time, drop
    it from our private state structure, and remove our use of crtc->state
    from our vblank handler entirely.

    Link: https://lore.kernel.org/all/YWgteNaNeaS9uWDe@phenom.ffwll.local/
    Link: https://lore.kernel.org/r/20211025141113.702757-4-maxime@cerno.tech
    Fixes: 87ebcd42fb7b ("drm/vc4: crtc: Assign output to channel automatically")
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c3eb3d168ac8049656d65c2f11d1f02e69639d05
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Oct 25 16:11:06 2021 +0200

    drm/vc4: Fix non-blocking commit getting stuck forever

    commit 0c250c150c74a90db298bf2a8bcd0a1dabed2e2f upstream.

    In some situation, we can end up being stuck on a non-blocking that went
    through properly.

    The situation that seems to trigger it reliably is to first start a
    non-blocking commit, and then right after, and before we had any vblank
    interrupt), start a blocking commit.

    This will lead to the first commit workqueue to be scheduled, setup the
    display, while the second commit is waiting for the first one to be
    completed.

    The vblank interrupt will then be raised, vc4_crtc_handle_vblank() will
    run and will compare the active dlist in the HVS channel to the one
    associated with the crtc->state.

    However, at that point, the second commit is waiting using
    drm_atomic_helper_wait_for_dependencies that occurs after
    drm_atomic_helper_swap_state has been called, so crtc->state points to
    the second commit state. vc4_crtc_handle_vblank() will compare the two
    dlist addresses and since they don't match will ignore the interrupt.

    The vblank event will never be reported, and the first and second commit
    will wait for the first commit completion until they timeout.

    The underlying reason is that it was never safe to do so. Indeed,
    accessing the ->state pointer access synchronization is based on
    ownership guarantees that can only occur within the functions and hooks
    defined as part of the KMS framework, and obviously the irq handler
    isn't one of them. The rework to move to generic helpers only uncovered
    the underlying issue.

    However, since the code path between
    drm_atomic_helper_wait_for_dependencies() and
    drm_atomic_helper_wait_for_vblanks() is serialised and we can't get two
    commits in that path at the same time, we can work around this issue by
    setting a variable associated to struct drm_crtc to the dlist we expect,
```

```
    and then using it from the vc4_crtc_handle_vblank() function.

    Since that state is shared with the modesetting path, we also need to
    introduce a spinlock to protect the code shared between the interrupt
    handler and the modesetting path, protecting only our new variable for
    now.

    Link: https://lore.kernel.org/all/YWgteNaNeaS9uWDe@phenom.ffwll.local/
    Link: https://lore.kernel.org/r/20211025141113.702757-3-maxime@cerno.tech
    Fixes: 56d1fe0979dc ("drm/vc4: Make pageflip completion handling more robust.")
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 40bc3d3b52471e9584180992fd7148f0d7813ce0
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Oct 25 16:11:05 2021 +0200

    drm/vc4: crtc: Drop feed_txp from state

    commit a16c66401fd831f70a02d33e9bcaac585637c29f upstream.

    Accessing the crtc->state pointer from outside the modesetting context
    is not allowed. We thus need to copy whatever we need from the KMS state
    to our structure in order to access it.

    In VC4, a number of users of that pointers have crept in over the years,
    the first one being whether or not the downstream controller of the
    pixelvalve is our writeback controller.

    Fortunately for us, Since commit 39fcb2808376 ("drm/vc4: txp: Turn the
    TXP into a CRTC of its own") this is no longer something that can change
    from one commit to the other and is hardcoded.

    Let's set this flag in struct vc4_crtc if we happen to be the TXP, and
    drop the flag from our private state structure.

    Link: https://lore.kernel.org/all/YWgteNaNeaS9uWDe@phenom.ffwll.local/
    Link: https://lore.kernel.org/r/20211025141113.702757-2-maxime@cerno.tech
    Fixes: 008095e065a8 ("drm/vc4: Add support for the transposer block")
    Acked-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 161e5bb0b471cf252855bcaf9acbaa331161c1b5
Author: Ye Bin <yebin10@huawei.com>
Date:   Mon Nov 29 09:26:59 2021 +0800

    block: Fix fsync always failed if once failed

    commit 8a7518931baa8ea023700987f3db31cb0a80610b upstream.

    We do test with inject error fault base on v4.19, after test some time we found
    sync /dev/sda always failed.
    [root@localhost] sync /dev/sda
    sync: error syncing '/dev/sda': Input/output error

    scsi log as follows:
    [19069.812296] sd 0:0:0:0: [sda] tag#64 Send: scmd 0x00000000d03a0b6b
    [19069.812302] sd 0:0:0:0: [sda] tag#64 CDB: Synchronize Cache(10) 35 00 00 00 00 00 00
00 00 00
    [19069.812533] sd 0:0:0:0: [sda] tag#64 Done: SUCCESS Result: hostbyte=DID_OK
driverbyte=DRIVER_OK
    [19069.812536] sd 0:0:0:0: [sda] tag#64 CDB: Synchronize Cache(10) 35 00 00 00 00 00 00
00 00 00
```

```
     [19069.812539] sd 0:0:0:0: [sda] tag#64 scsi host busy 1 failed 0
     [19069.812542] sd 0:0:0:0: Notifying upper driver of completion (result 0)
     [19069.812546] sd 0:0:0:0: [sda] tag#64 sd_done: completed 0 of 0 bytes
     [19069.812549] sd 0:0:0:0: [sda] tag#64 0 sectors total, 0 bytes done.
     [19069.812564] print_req_error: I/O error, dev sda, sector 0

     ftrace log as follows:
      rep-306069 [007] .... 19654.923315: block_bio_queue: 8,0 FWS 0 + 0 [rep]
      rep-306069 [007] .... 19654.923333: block_getrq: 8,0 FWS 0 + 0 [rep]
      kworker/7:1H-250   [007] .... 19654.923352: block_rq_issue: 8,0 FF 0 () 0 + 0
 [kworker/7:1H]
      <idle>-0     [007] ..s. 19654.923562: block_rq_complete: 8,0 FF () 18446744073709551615
 + 0 [0]
      <idle>-0     [007] d.s. 19654.923576: block_rq_complete: 8,0 WS () 0 + 0 [-5]

     As 8d6996630c03 introduce 'fq->rq_status', this data only update when 'flush_rq'
     reference count isn't zero. If flush request once failed and record error code
     in 'fq->rq_status'. If there is no chance to update 'fq->rq_status',then do fsync
     will always failed.
     To address this issue reset 'fq->rq_status' after return error code to upper layer.

     Fixes: 8d6996630c03("block: fix null pointer dereference in blk_mq_rq_timed_out()")
     Signed-off-by: Ye Bin <yebin10@huawei.com>
     Reviewed-by: Ming Lei <ming.lei@redhat.com>
     Link: https://lore.kernel.org/r/20211129012659.1553733-1-yebin10@huawei.com
     Signed-off-by: Jens Axboe <axboe@kernel.dk>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 9d5449c54207d7ac2406396e0b2fe7016ae88183
 Author: Jens Axboe <axboe@kernel.dk>
 Date:   Thu Jan 20 10:28:13 2022 -0700

     block: fix async_depth sysfs interface for mq-deadline

     commit 46cdc45acb089c811d9a54fd50af33b96e5fae9d upstream.

     A previous commit added this feature, but it inadvertently used the wrong
     variable to show/store the setting from/to, victimized by copy/paste. Fix
     it up so that the async_depth sysfs interface reads and writes from the
     right setting.

     Fixes: 07757588e507 ("block/mq-deadline: Reserve 25% of scheduler tags for synchronous
 requests")
     Link: https://bugzilla.kernel.org/show_bug.cgi?id=215485
     Reviewed-by: Bart Van Assche <bvanassche@acm.org>
     Signed-off-by: Jens Axboe <axboe@kernel.dk>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 521d64a9912fbff2e449d645b35b520485dab486
 Author: Tobias Waldekranz <tobias@waldekranz.com>
 Date:   Tue Jan 18 22:50:52 2022 +0100

     powerpc/fsl/dts: Enable WA for erratum A-009885 on fman3l MDIO buses

     commit 0d375d610fa96524e2ee2b46830a46a7bfa92a9f upstream.

     This block is used in (at least) T1024 and T1040, including their
     variants like T1023 etc.

     Fixes: d55ad2967d89 ("powerpc/mpc85xx: Create dts components for the FSL QorIQ DPAA
 FMan")
     Signed-off-by: Tobias Waldekranz <tobias@waldekranz.com>
     Signed-off-by: Jakub Kicinski <kuba@kernel.org>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit cf93f14958c029d2774b43e5e5d381647e84067f
Author: Anders Roxell <anders.roxell@linaro.org>
Date:   Tue Dec 7 12:02:28 2021 +0100

    powerpc/cell: Fix clang -Wimplicit-fallthrough warning

    commit e89257e28e844f5d1d39081bb901d9f1183a7705 upstream.

    Clang warns:

    arch/powerpc/platforms/cell/pervasive.c:81:2: error: unannotated fall-through between
switch labels
            case SRR1_WAKEEE:
            ^
    arch/powerpc/platforms/cell/pervasive.c:81:2: note: insert 'break;' to avoid fall-
through
            case SRR1_WAKEEE:
            ^
            break;
    1 error generated.

    Clang is more pedantic than GCC, which does not warn when failing
    through to a case that is just break or return. Clang's version is more
    in line with the kernel's own stance in deprecated.rst. Add athe missing
    break to silence the warning.

    Fixes: 6e83985b0f6e ("powerpc/cbe: Do not process external or decremeter interrupts from
sreset")
    Reported-by: Naresh Kamboju <naresh.kamboju@linaro.org>
    Signed-off-by: Anders Roxell <anders.roxell@linaro.org>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211207110228.698956-1-anders.roxell@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ed267a535cba22b9f4fa59b269d428b92d4dfdf1
Author: Moshe Shemesh <moshe@nvidia.com>
Date:   Sun Dec 5 11:20:59 2021 +0200

    Revert "net/mlx5: Add retry mechanism to the command entry index allocation"

    commit 4f6626b0e140867fd6d5a2e9d4ceaef97f10f46a upstream.

    This reverts commit 410bd754cd73c4a2ac3856d9a03d7b08f9c906bf.

    The reverted commit had added a retry mechanism to the command entry
    index allocation. The previous patch ensures that there is a free
    command entry index once the command work handler holds the command
    semaphore. Thus the retry mechanism is not needed.

    Fixes: 410bd754cd73 ("net/mlx5: Add retry mechanism to the command entry index
allocation")
    Signed-off-by: Moshe Shemesh <moshe@nvidia.com>
    Reviewed-by: Eran Ben Elisha <eranbe@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b6d51ea8cbb351efa1634d96944d21ac72052df1
Author: Amelie Delaunay <amelie.delaunay@foss.st.com>
Date:   Mon Dec 20 17:58:27 2021 +0100

    dmaengine: stm32-mdma: fix STM32_MDMA_CTBR_TSEL_MASK

    commit e7f110889a87307fb0fed408a5dee1707796ca04 upstream.
```

```
    This patch fixes STM32_MDMA_CTBR_TSEL_MASK, which is [5:0], not [7:0].

    Fixes: a4ffb13c8946 ("dmaengine: Add STM32 MDMA driver")
    Signed-off-by: Amelie Delaunay <amelie.delaunay@foss.st.com>
    Link: https://lore.kernel.org/r/20211220165827.1238097-1-amelie.delaunay@foss.st.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d6527cdee1b5f6b74745ec9be768ef4ff12b952a
Author: Chengguang Xu <cgxu519@mykernel.net>
Date:   Sat Dec 18 19:23:20 2021 +0800

    RDMA/rxe: Fix a typo in opcode name

    commit 8d1cfb884e881efd69a3be4ef10772c71cb22216 upstream.

    There is a redundant ']' in the name of opcode IB_OPCODE_RC_SEND_MIDDLE,
    so just fix it.

    Fixes: 8700e3e7c485 ("Soft RoCE driver")
    Link: https://lore.kernel.org/r/20211218112320.3558770-1-cgxu519@mykernel.net
    Signed-off-by: Chengguang Xu <cgxu519@mykernel.net>
    Acked-by: Zhu Yanjun <zyjzyj2000@gmail.com>
    Reviewed-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e68b71737c81c90849cd51734fe0b54d4d1ac182
Author: Yixing Liu <liuyixing1@huawei.com>
Date:   Mon Dec 6 21:36:52 2021 +0800

    RDMA/hns: Modify the mapping attribute of doorbell to device

    commit 39d5534b1302189c809e90641ffae8cbdc42a8fc upstream.

    It is more general for ARM device drivers to use the device attribute to
    map PCI BAR spaces.

    Fixes: 9a4435375cd1 ("IB/hns: Add driver files for hns RoCE driver")
    Link: https://lore.kernel.org/r/20211206133652.27476-1-liangwenpeng@huawei.com
    Signed-off-by: Yixing Liu <liuyixing1@huawei.com>
    Signed-off-by: Wenpeng Liang <liangwenpeng@huawei.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d811d2531edc98f84f42098ed841a44f2fce29c5
Author: Dave Jiang <dave.jiang@intel.com>
Date:   Tue Dec 14 13:15:17 2021 -0700

    dmaengine: idxd: fix wq settings post wq disable

    commit 0f225705cf6536826318180831e18a74595efc8d upstream.

    By the spec, wq size and group association is not changeable unless device
    is disabled. Exclude clearing the shadow copy on wq disable/reset. This
    allows wq type to be changed after disable to be re-enabled.

    Move the size and group association to its own cleanup and only call it
    during device disable.

    Fixes: 0dcfe41e9a4c ("dmanegine: idxd: cleanup all device related bits after disabling
device")
    Reported-by: Lucas Van <lucas.van@intel.com>
    Tested-by: Lucas Van <lucas.van@intel.com>
```

```
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/163951291732.2987775.13576571320501115257.stgit@djiang5-
desk3.ch.intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8e873a76a6b17cbc23024510374286a8d27523a1
Author: Kunihiko Hayashi <hayashi.kunihiko@socionext.com>
Date:   Tue Dec 14 13:42:43 2021 +0900

    dmaengine: uniphier-xdmac: Fix type of address variables

    commit 105a8c525675bb7d4d64871f9b2edf39460de881 upstream.

    The variables src_addr and dst_addr handle DMA addresses, so these should
    be declared as dma_addr_t.

    Fixes: 667b9251440b ("dmaengine: uniphier-xdmac: Add UniPhier external DMA controller
driver")
    Signed-off-by: Kunihiko Hayashi <hayashi.kunihiko@socionext.com>
    Link: https://lore.kernel.org/r/1639456963-10232-1-git-send-email-
hayashi.kunihiko@socionext.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c994dbcc58d6d09c0f736dfbf9f5d6a62d26bd5a
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Wed Dec 22 07:09:30 2021 +0000

    scsi: ufs: ufs-mediatek: Fix error checking in ufs_mtk_init_va09_pwr_ctrl()

    commit 3ba880a12df5aa4488c18281701b5b1bc3d4531a upstream.

    The function regulator_get() returns an error pointer. Use IS_ERR() to
    validate the return value.

    Link: https://lore.kernel.org/r/20211222070930.9449-1-linmq006@gmail.com
    Fixes: cf137b3ea49a ("scsi: ufs-mediatek: Support VA09 regulator operations")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9b57e75bc30e1b74aa4e226d191235f166d84997
Author: Baruch Siach <baruch@tkos.co.il>
Date:   Thu Dec 30 18:31:53 2021 +0200

    of: base: Improve argument length mismatch error

    commit 5d05b811b5acb92fc581a7b328b36646c86f5ab9 upstream.

    The cells_name field of of_phandle_iterator might be NULL. Use the
    phandle name instead. With this change instead of:

      OF: /soc/pinctrl@1000000: (null) = 3 found 2

    We get:

      OF: /soc/pinctrl@1000000: phandle pinctrl@1000000 needs 3, found 2

    Which is a more helpful messages making DT debugging easier.

    In this particular example the phandle name looks like duplicate of the
    same node name. But note that the first node is the parent node
    (it->parent), while the second is the phandle target (it->node). They
    happen to be the same in the case that triggered this improvement. See
```

```
      commit 72cb4c48a46a ("arm64: dts: qcom: ipq6018: Fix gpio-ranges
      property").

      Signed-off-by: Baruch Siach <baruch@tkos.co.il>
      Signed-off-by: Rob Herring <robh@kernel.org>
      Link:
  https://lore.kernel.org/r/f6a68e0088a552ea9dfd4d8e3b5b586d92594738.1640881913.git.baruch@tko
  s.co.il
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit ff2d9c21832ba69a3043d87ae38b8ef7c2c86314
  Author: Bart Van Assche <bvanassche@acm.org>
  Date:   Mon Nov 29 11:46:00 2021 -0800

      scsi: core: Show SCMD_LAST in text form

      commit 3369046e54ca8f82e0cb17740643da2d80d3cfa8 upstream.

      The SCSI debugfs code supports showing information about pending commands,
      including translating SCSI command flags from numeric into text format.
      Also convert the SCMD_LAST flag from numeric into text form.

      Link: https://lore.kernel.org/r/20211129194609.3466071-4-bvanassche@acm.org
      Fixes: 8930a6c20791 ("scsi: core: add support for request batching")
      Signed-off-by: Bart Van Assche <bvanassche@acm.org>
      Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 5e51917862d8075889eb7afd6edb115209e21e90
  Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
  Date:   Mon Nov 1 15:14:41 2021 -0700

      Bluetooth: hci_sync: Fix not setting adv set duration

      commit f16a491c65d9eb19398b25aefc10c2d3313d17b3 upstream.

      10bbffa3e88e attempted to fix the use of rotation duration as
      advertising duration but it didn't change the if condition which still
      uses the duration instead of the timeout.

      Fixes: 10bbffa3e88e ("Bluetooth: Fix using advertising instance duration as timeout")
      Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
      Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 8cb8d0b940ba813f251a51d40ae583d0a39e3c44
  Author: Markus Reichl <m.reichl@fivetechno.de>
  Date:   Thu Jan 13 21:01:11 2022 +0100

      net: usb: Correct reset handling of smsc95xx

      commit 0bf3885324a8599e3af4c7379b8d4f621c9bbffa upstream.

      On boards with LAN9514 and no preconfigured MAC address we don't get an
      ip address from DHCP after commit a049a30fc27c ("net: usb: Correct PHY handling
      of smsc95xx") anymore. Adding an explicit reset before starting the phy
      fixes the issue.

      [1]

  https://lore.kernel.org/netdev/199eebbd6b97f52b9119c9fa4fd8504f8a34de18.camel@collabora.com/

      From: Gabriel Hojda <ghojda@yo2urs.ro>
      Fixes: a049a30fc27c ("net: usb: Correct PHY handling of smsc95xx")
      Signed-off-by: Gabriel Hojda <ghojda@yo2urs.ro>
```

```
    Signed-off-by: Markus Reichl <m.reichl@fivetechno.de>
    Tested-by: Alexander Stein <alexander.stein@ew.tq-group.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 29d29f8970b4751a8e6fbdc3ebbaf48ed1496dec
Author: Mark Chen <mark-yw.chen@mediatek.com>
Date:   Tue Dec 7 01:33:43 2021 +0800

    Bluetooth: btusb: Return error code when getting patch status failed

    commit 995d948cf2e45834275f07afc1c9881a9902e73c upstream.

    If there are failure cases in getting patch status, it should return the
    error code (-EIO).

    Fixes: fc342c4dc4087 ("Bluetooth: btusb: Add protocol support for MediaTek MT7921U USB
devices")
    Co-developed-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Mark Chen <mark-yw.chen@mediatek.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 055f208e408a3bb05634c99a18c935d2ed6dbddc
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Tue Jan 18 19:39:05 2022 -0800

    Documentation: fix firewire.rst ABI file path error

    commit b0ac702f3329cdc8a06dcaac73183d4b5a2b942d upstream.

    Adjust the path of the ABI files for firewire.rst to prevent a
    documentation build error. Prevents this problem:

    Sphinx parallel build error:
    docutils.utils.SystemMessage: Documentation/driver-api/firewire.rst:22: (SEVERE/4)
Problems with "include" directive path:
    InputError: [Errno 2] No such file or directory: '../Documentation/driver-
api/ABI/stable/firewire-cdev'.

    Fixes: 2f4830ef96d2 ("FireWire: add driver-api Introduction section")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Tested-by: Akira Yokosawa <akiyks@gmail.com>
    Link: https://lore.kernel.org/r/20220119033905.4779-1-rdunlap@infradead.org
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1d6b92f3ae5fc1a5dd28915ce73bd97a27b54181
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Thu Dec 30 18:19:40 2021 +0100

    Documentation: refer to config RANDOMIZE_BASE for kernel address-space randomization

    commit 82ca67321f55a8d1da6ac3ed611da3c32818bb37 upstream.

    The config RANDOMIZE_SLAB does not exist, the authors probably intended to
    refer to the config RANDOMIZE_BASE, which provides kernel address-space
    randomization. They probably just confused SLAB with BASE (these two
    four-letter words coincidentally share three common letters), as they also
    point out the config SLAB_FREELIST_RANDOM as further randomization within
    the same sentence.

    Fix the reference of the config for kernel address-space randomization to
    the config that provides that.
```

```
    Fixes: 6e88559470f5 ("Documentation: Add section about CPU vulnerabilities for Spectre")
    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Link: https://lore.kernel.org/r/20211230171940.27558-1-lukas.bulwahn@gmail.com
    Signed-off-by: Jonathan Corbet <corbet@lwn.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 33150ae6236f56f64bd4c85dee5a8de789395eb5
Author: Alexandre Ghiti <alexandre.ghiti@canonical.com>
Date:   Thu Dec 16 10:44:23 2021 +0100

    Documentation, arch: Remove leftovers from CIFS_WEAK_PW_HASH

    commit 2ac7069ad7647cd1d9ca5b08765a1e116e13cdc4 upstream.

    This config was removed so remove all references to it.

    Fixes: 76a3c92ec9e0 ("cifs: remove support for NTLM and weaker authentication
algorithms")
    Signed-off-by: Alexandre Ghiti <alexandre.ghiti@canonical.com>
    Reviewed-by: Steve French <smfrench@gmail.com>
    Acked-by: Arnd Bergmann <arnd@arndb.de> [arch/arm/configs]
    Acked-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3bee8b34dd8e487654322b682c1580acea27a5ba
Author: Alexandre Ghiti <alexandre.ghiti@canonical.com>
Date:   Thu Dec 16 10:44:22 2021 +0100

    Documentation, arch: Remove leftovers from raw device

    commit 473dcf0ffc31ce1135cd10578e7e06698cf51f4a upstream.

    Raw device interface was removed so remove all references to configs
    related to it.

    Fixes: 603e4922f1c8 ("remove the raw driver")
    Signed-off-by: Alexandre Ghiti <alexandre.ghiti@canonical.com>
    Acked-by: Arnd Bergmann <arnd@arndb.de> [arch/arm/configs]
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f33c24e34ddcda0fdacd6a0b3bcc2ae84b3aff8b
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Wed Dec 1 14:59:31 2021 +0200

    Documentation: ACPI: Fix data node reference documentation

    commit a11174952205d082f1658fab4314f0caf706e0a8 upstream.

    The data node reference documentation was missing a package that must
    contain the property values, instead property name and multiple values
    being present in a single package. This is not aligned with the _DSD
    spec.

    Fix it by adding the package for the values.

    Also add the missing "reg" properties to two numbered nodes.

    Fixes: b10134a3643d ("ACPI: property: Document hierarchical data extension references")
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 75a766008e5cbc78cb9c6005d60e177e7d0ae146
Author: Daniel Thompson <daniel.thompson@linaro.org>
Date:   Thu Nov 18 10:09:52 2021 +0000

    Documentation: dmaengine: Correctly describe dmatest with channel unset

    commit c61d7b2ef141abf81140756b45860a2306f395a2 upstream.

    Currently the documentation states that channels must be configured before
    running the dmatest. This has not been true since commit 6b41030fdc79
    ("dmaengine: dmatest: Restore default for channel"). Fix accordingly.

    Fixes: 6b41030fdc79 ("dmaengine: dmatest: Restore default for channel")
    Signed-off-by: Daniel Thompson <daniel.thompson@linaro.org>
    Link: https://lore.kernel.org/r/20211118100952.27268-3-daniel.thompson@linaro.org
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8fb5dd311d5711cf9ca5f432e90372b00def1f36
Author: Mike Leach <mike.leach@linaro.org>
Date:   Wed Nov 17 16:42:20 2021 +0000

    Documentation: coresight: Fix documentation issue

    commit 66bd1333abd7fa191f13b929c9119d6cd3df27b0 upstream.

    Fix the description of the directories and attributes used
    in cs_etm as used by perf.

    Drop the references to the 'configurations' sub-directory which
    had been removed in an earlier version of the patchset.

    Fixes: f71cd93d5ea4 ("Documentation: coresight: Add documentation for CoreSight config")
    Reported-by: German Gomex <german.gomez@arm.com>
    Signed-off-by: Mike Leach <mike.leach@linaro.org>
    Link: https://lore.kernel.org/r/20211117164220.14883-1-mike.leach@linaro.org
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4b55b553508be3cd9c1c1556cd0a2f91c14434a8
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Nov 7 18:19:23 2021 +0000

    media: correct MEDIA_TEST_SUPPORT help text

    commit 09f4d1513267d0ab712f5d29e7bd136535748709 upstream.

    Fix grammar/wording in the help text for MEDIA_TEST_SUPPORT.

    Fixes: 4b32216adb01 ("media: split test drivers from platform directory")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.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 9b0d360fd783c711fc1cafa51f3e03bdf8ca5518
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Thu Aug 19 15:59:30 2021 +0200

    drm/vc4: hdmi: Make sure the device is powered with CEC

    commit 20b0dfa86bef0e80b41b0e5ac38b92f23b6f27f9 upstream.

    Similarly to what we encountered with the detect hook with DRM, nothing
```

```
    actually prevents any of the CEC callback from being run while the HDMI
    output is disabled.

    However, this is an issue since any register access to the controller
    when it's powered down will result in a silent hang.

    Let's make sure we run the runtime_pm hooks when the CEC adapter is
    opened and closed by the userspace to avoid that issue.

    Fixes: 15b4511a4af6 ("drm/vc4: add HDMI CEC support")
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210819135931.895976-6-
maxime@cerno.tech
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e3032bd03e5471c97eca6a59eeb045e2935c26d0
Author: Suresh Udipi <sudipi@jp.adit-jv.com>
Date:   Fri Aug 13 17:07:56 2021 +0200

    media: rcar-csi2: Optimize the selection PHTW register

    commit 549cc89cd09a85aaa16dc07ef3db811d5cf9bcb1 upstream.

    PHTW register is selected based on default bit rate from Table[1].
    for the bit rates less than or equal to 250. Currently first
    value of default bit rate which is greater than or equal to
    the caculated mbps is selected. This selection can be further
    improved by selecting the default bit rate which is nearest to
    the calculated value.

    [1] specs r19uh0105ej0200-r-car-3rd-generation.pdf [Table 25.12]

    Fixes: 769afd212b16 ("media: rcar-csi2: add Renesas R-Car MIPI CSI-2 receiver driver")
    Signed-off-by: Suresh Udipi <sudipi@jp.adit-jv.com>
    Signed-off-by: Michael Rodin <mrodin@de.adit-jv.com>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    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 69134ac6836213b0012d7a50485eda0a33844beb
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Fri Apr 23 11:26:56 2021 +0200

    can: mcp251xfd: mcp251xfd_tef_obj_read(): fix typo in error message

    commit 99e7cc3b3f85d9a583ab83f386315c59443509ae upstream.

    This patch fixes a typo in the error message in
    mcp251xfd_tef_obj_read(), if trying to read too many objects.

    Link: https://lore.kernel.org/all/20220105154300.1258636-3-mkl@pengutronix.de
    Fixes: 55e5b97f003e ("can: mcp25xxfd: add driver for Microchip MCP25xxFD SPI CAN")
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 220aacbd122a3a738cf4b53e91968042307d47db
Author: Ben Hutchings <ben@decadent.org.uk>
Date:   Mon Jun 18 23:55:40 2018 +0100

    firmware: Update Kconfig help text for Google firmware

    commit d185a3466f0cd5af8f1c5c782c53bc0e6f2e7136 upstream.
```

```
    The help text for GOOGLE_FIRMWARE states that it should only be
    enabled when building a kernel for Google's own servers.  However,
    many of the drivers dependent on it are also useful on Chromebooks or
    on any platform using coreboot.

    Update the help text to reflect this double duty.

    Fixes: d384d6f43d1e ("firmware: google memconsole: Add coreboot support")
    Reviewed-by: Julius Werner <jwerner@chromium.org>
    Signed-off-by: Ben Hutchings <ben@decadent.org.uk>
    Link: https://lore.kernel.org/r/20180618225540.GD14131@decadent.org.uk
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9b510908e42084820c9d71abf0b808bda2315dde
Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
Date:   Tue Nov 30 09:32:33 2021 -0500

    drm/amdgpu/display: Only set vblank_disable_immediate when PSR is not enabled

    commit 70897848730470cc477d5d89e6222c0f6a9ac173 upstream.

    [Why]
    PSR currently relies on the kernel's delayed vblank on/off mechanism
    as an implicit bufferring mechanism to prevent excessive entry/exit.

    Without this delay the user experience is impacted since it can take
    a few frames to enter/exit.

    [How]
    Only allow vblank disable immediate for DC when psr is not supported.

    Leave a TODO indicating that this support should be extended in the
    future to delay independent of the vblank interrupt.

    Fixes: 92020e81ddbeac ("drm/amdgpu/display: set vblank_disable_immediate for DC")

    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f2d67e7000d11a9dbba1ef3d612fe001601c6eb4
Author: Christian König <christian.koenig@amd.com>
Date:   Mon Jan 17 10:31:26 2022 +0100

    drm/radeon: fix error handling in radeon_driver_open_kms

    commit 4722f463896cc0ef1a6f1c3cb2e171e949831249 upstream.

    The return value was never initialized so the cleanup code executed when
    it isn't even necessary.

    Just add proper error handling.

    Fixes: ab50cb9df889 ("drm/radeon/radeon_kms: Fix a NULL pointer dereference in
radeon_driver_open_kms()")
    Signed-off-by: Christian König <christian.koenig@amd.com>
    Tested-by: Jan Stancek <jstancek@redhat.com>
    Tested-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 09589fac2cb4da9ba14bf65d8daa0f846a37bf22
```

```
Author: Theodore Ts'o <tytso@mit.edu>
Date:   Wed Jan 5 23:59:56 2022 -0500

    ext4: don't use the orphan list when migrating an inode

    commit 6eeaf88fd586f05aaf1d48cb3a139d2a5c6eb055 upstream.

    We probably want to remove the indirect block to extents migration
    feature after a deprecation window, but until then, let's fix a
    potential data loss problem caused by the fact that we put the
    tmp_inode on the orphan list.  In the unlikely case where we crash and
    do a journal recovery, the data blocks belonging to the inode being
    migrated are also represented in the tmp_inode on the orphan list ---
    and so its data blocks will get marked unallocated, and available for
    reuse.

    Instead, stop putting the tmp_inode on the oprhan list.  So in the
    case where we crash while migrating the inode, we'll leak an inode,
    which is not a disaster.  It will be easily fixed the next time we run
    fsck, and it's better than potentially having blocks getting claimed
    by two different files, and losing data as a result.

    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Reviewed-by: Lukas Czerner <lczerner@redhat.com>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f9152b7fe640c73ec0efa229c811a123871a7b65
Author: Zhang Yi <yi.zhang@huawei.com>
Date:   Sat Dec 25 17:09:37 2021 +0800

    ext4: fix an use-after-free issue about data=journal writeback mode

    commit 5c48a7df91499e371ef725895b2e2d21a126e227 upstream.

    Our syzkaller report an use-after-free issue that accessing the freed
    buffer_head on the writeback page in __ext4_journalled_writepage(). The
    problem is that if there was a truncate racing with the data=journalled
    writeback procedure, the writeback length could become zero and
    bget_one() refuse to get buffer_head's refcount, then the truncate
    procedure release buffer once we drop page lock, finally, the last
    ext4_walk_page_buffers() trigger the use-after-free problem.

    sync                                    truncate
    ext4_sync_file()
     file_write_and_wait_range()
                                            ext4_setattr(0)
                                             inode->i_size = 0
      ext4_writepage()
       len = 0
       __ext4_journalled_writepage()
        page_bufs = page_buffers(page)
        ext4_walk_page_buffers(bget_one) <- does not get refcount
                                             do_invalidatepage()
                                               free_buffer_head()
        ext4_walk_page_buffers(page_bufs) <- trigger use-after-free

    After commit bdf96838aea6 ("ext4: fix race between truncate and
    __ext4_journalled_writepage()"), we have already handled the racing
    case, so the bget_one() and bput_one() are not needed. So this patch
    simply remove these hunk, and recheck the i_size to make it safe.

    Fixes: bdf96838aea6 ("ext4: fix race between truncate and
__ext4_journalled_writepage()")
    Signed-off-by: Zhang Yi <yi.zhang@huawei.com>
```

```
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20211225090937.712867-1-yi.zhang@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit acc2e0a1a6c5493f0b1be95a7cc489a77291ce14
Author: Ye Bin <yebin10@huawei.com>
Date:   Fri Dec 24 18:03:41 2021 +0800

    ext4: fix null-ptr-deref in '__ext4_journal_ensure_credits'

    commit 298b5c521746d69c07beb2757292fb5ccc1b0f85 upstream.

    We got issue as follows when run syzkaller test:
    [ 1901.130043] EXT4-fs error (device vda): ext4_remount:5624: comm syz-executor.5: Abort
forced by user
    [ 1901.130901] Aborting journal on device vda-8.
    [ 1901.131437] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.16: Detected aborted journal
    [ 1901.131566] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.11: Detected aborted journal
    [ 1901.132586] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.18: Detected aborted journal
    [ 1901.132751] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.9: Detected aborted journal
    [ 1901.136149] EXT4-fs error (device vda) in ext4_reserve_inode_write:6035: Journal has
aborted
    [ 1901.136837] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-fuzzer:
Detected aborted journal
    [ 1901.136915] ================================================================
    [ 1901.138175] BUG: KASAN: null-ptr-deref in __ext4_journal_ensure_credits+0x74/0x140
[ext4]
    [ 1901.138343] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.13: Detected aborted journal
    [ 1901.138398] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.1: Detected aborted journal
    [ 1901.138808] Read of size 8 at addr 0000000000000000 by task syz-executor.17/968
    [ 1901.138817]
    [ 1901.138852] EXT4-fs error (device vda): ext4_journal_check_start:61: comm syz-
executor.30: Detected aborted journal
    [ 1901.144779] CPU: 1 PID: 968 Comm: syz-executor.17 Not tainted 4.19.90-
vhulk2111.1.0.h893.eulerosv2r10.aarch64+ #1
    [ 1901.146479] Hardware name: linux,dummy-virt (DT)
    [ 1901.147317] Call trace:
    [ 1901.147552]  dump_backtrace+0x0/0x2d8
    [ 1901.147898]  show_stack+0x28/0x38
    [ 1901.148215]  dump_stack+0xec/0x15c
    [ 1901.148746]  kasan_report+0x108/0x338
    [ 1901.149207]  __asan_load8+0x58/0xb0
    [ 1901.149753]  __ext4_journal_ensure_credits+0x74/0x140 [ext4]
    [ 1901.150579]  ext4_xattr_delete_inode+0xe4/0x700 [ext4]
    [ 1901.151316]  ext4_evict_inode+0x524/0xba8 [ext4]
    [ 1901.151985]  evict+0x1a4/0x378
    [ 1901.152353]  iput+0x310/0x428
    [ 1901.152733]  do_unlinkat+0x260/0x428
    [ 1901.153056]  __arm64_sys_unlinkat+0x6c/0xc0
    [ 1901.153455]  el0_svc_common+0xc8/0x320
    [ 1901.153799]  el0_svc_handler+0xf8/0x160
    [ 1901.154265]  el0_svc+0x10/0x218
    [ 1901.154682] ================================================================

    This issue may happens like this:
            Process1                                Process2
    ext4_evict_inode
      ext4_journal_start
```

```
        ext4_truncate
          ext4_ind_truncate
            ext4_free_branches
              ext4_ind_truncate_ensure_credits
                ext4_journal_ensure_credits_fn
                  ext4_journal_restart
                    handle->h_transaction = NULL;
                                              mount -o remount,abort  /mnt
                                              -> trigger JBD abort
                    start_this_handle -> will return failed
       ext4_xattr_delete_inode
         ext4_journal_ensure_credits
           ext4_journal_ensure_credits_fn
             __ext4_journal_ensure_credits
               jbd2_handle_buffer_credits
                 journal = handle->h_transaction->t_journal; ->null-ptr-deref
```

    Now, indirect truncate process didn't handle error. To solve this issue
    maybe simply add check handle is abort in '__ext4_journal_ensure_credits'
    is enough, and i also think this is necessary.

    Cc: stable@kernel.org
    Signed-off-by: Ye Bin <yebin10@huawei.com>
    Link: https://lore.kernel.org/r/20211224100341.3299128-1-yebin10@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f2cb343e44328a0758ff765a2088dab8a236fce5
Author: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
Date: Thu Dec 23 17:44:36 2021 +0100

    ext4: destroy ext4_fc_dentry_cachep kmemcache on module removal

    commit ab047d516dea72f011c15c04a929851e4d053109 upstream.

    The kmemcache for ext4_fc_dentry_cachep remains registered after module
    removal.

    Destroy ext4_fc_dentry_cachep kmemcache on module removal.

    Fixes: aa75f4d3daaeb ("ext4: main fast-commit commit path")
    Signed-off-by: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
    Reviewed-by: Lukas Czerner <lczerner@redhat.com>
    Reviewed-by: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
    Link: https://lore.kernel.org/r/20211110134640.lyku5vklvdndw6uk@linutronix.de
    Link: https://lore.kernel.org/r/YbiK3JetFFl08bd7@linutronix.de
    Link: https://lore.kernel.org/r/20211223164436.2628390-1-bigeasy@linutronix.de
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e475b3350c973e8aea45caa7ca7c8a6369a86383
Author: Xin Yin <yinxin.x@bytedance.com>
Date: Thu Dec 23 11:23:37 2021 +0800

    ext4: fast commit may miss tracking unwritten range during ftruncate

    commit 9725958bb75cdfa10f2ec11526fdb23e7485e8e4 upstream.

    If use FALLOC_FL_KEEP_SIZE to alloc unwritten range at bottom, the
    inode->i_size will not include the unwritten range. When call
    ftruncate with fast commit enabled, it will miss to track the
    unwritten range.

    Change to trace the full range during ftruncate.

```
    Signed-off-by: Xin Yin <yinxin.x@bytedance.com>
    Reviewed-by: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
    Link: https://lore.kernel.org/r/20211223032337.5198-3-yinxin.x@bytedance.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e92682c0c4881717d8e96c2363398e0d41b49866
Author: Xin Yin <yinxin.x@bytedance.com>
Date:   Thu Dec 23 11:23:36 2021 +0800

    ext4: use ext4_ext_remove_space() for fast commit replay delete range

    commit 0b5b5a62b945a141e64011b2f90ee7e46f14be98 upstream.

    For now ,we use ext4_punch_hole() during fast commit replay delete range
    procedure. But it will be affected by inode->i_size, which may not
    correct during fast commit replay procedure. The following test will
    failed.

    -create & write foo (len 1000K)
    -falloc FALLOC_FL_ZERO_RANGE foo (range 400K - 600K)
    -create & fsync bar
    -falloc FALLOC_FL_PUNCH_HOLE foo (range 300K-500K)
    -fsync foo
    -crash before a full commit

    After the fast_commit reply procedure, the range 400K-500K will not be
    removed. Because in this case, when calling ext4_punch_hole() the
    inode->i_size is 0, and it just retruns with doing nothing.

    Change to use ext4_ext_remove_space() instead of ext4_punch_hole()
    to remove blocks of inode directly.

    Signed-off-by: Xin Yin <yinxin.x@bytedance.com>
    Reviewed-by: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
    Link: https://lore.kernel.org/r/20211223032337.5198-2-yinxin.x@bytedance.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2997ab63284af87df88fbf71209445901be834bc
Author: Ye Bin <yebin10@huawei.com>
Date:   Thu Dec 23 09:55:06 2021 +0800

    ext4: Fix BUG_ON in ext4_bread when write quota data

    commit 380a0091cab482489e9b19e07f2a166ad2b76d5c upstream.

    We got issue as follows when run syzkaller:
    [  167.936972] EXT4-fs error (device loop0): __ext4_remount:6314: comm rep: Abort forced
by user
    [  167.938306] EXT4-fs (loop0): Remounting filesystem read-only
    [  167.981637] Assertion failure in ext4_getblk() at fs/ext4/inode.c:847:
'(EXT4_SB(inode->i_sb)->s_mount_state & EXT4_FC_REPLAY) || handle != NULL || create == 0'
    [  167.983601] ------------[ cut here ]------------
    [  167.984245] kernel BUG at fs/ext4/inode.c:847!
    [  167.984882] invalid opcode: 0000 [#1] PREEMPT SMP KASAN PTI
    [  167.985624] CPU: 7 PID: 2290 Comm: rep Tainted: G    B             5.16.0-rc5-next-
20211217+ #123
    [  167.986823] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS ?
-20190727_073836-buildvm-ppc64le-16.ppc.fedoraproject.org-3.fc31 04/01/2014
    [  167.988590] RIP: 0010:ext4_getblk+0x17e/0x504
    [  167.989189] Code: c6 01 74 28 49 c7 c0 a0 a3 5c 9b b9 4f 03 00 00 48 c7 c2 80 9c 5c
```

```
 9b 48 c7 c6 40 b6 5c 9b 48 c7 c7 20 a4 5c 9b e8 77 e3 fd ff <0f> 0b 8b 04 244
    [  167.991679] RSP: 0018:ffff8881736f7398 EFLAGS: 00010282
    [  167.992385] RAX: 0000000000000094 RBX: 1ffff1102e6dee75 RCX: 0000000000000000
    [  167.993337] RDX: 0000000000000001 RSI: ffffffff9b6e29e0 RDI: ffffed102e6dee66
    [  167.994292] RBP: ffff88816a076210 R08: 0000000000000094 R09: ffffed107363fa09
    [  167.995252] R10: ffff88839b1fd047 R11: ffffed107363fa08 R12: ffff88816a0761e8
    [  167.996205] R13: 0000000000000000 R14: 0000000000000021 R15: 0000000000000001
    [  167.997158] FS:  00007f6a1428c740(0000) GS:ffff88839b000000(0000)
 knlGS:0000000000000000
    [  167.998238] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [  167.999025] CR2: 00007f6a140716c8 CR3: 0000000133216000 CR4: 00000000000006e0
    [  167.999987] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [  168.000944] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [  168.001899] Call Trace:
    [  168.002235]  <TASK>
    [  168.007167]  ext4_bread+0xd/0x53
    [  168.007612]  ext4_quota_write+0x20c/0x5c0
    [  168.010457]  write_blk+0x100/0x220
    [  168.010944]  remove_free_dqentry+0x1c6/0x440
    [  168.011525]  free_dqentry.isra.0+0x565/0x830
    [  168.012133]  remove_tree+0x318/0x6d0
    [  168.014744]  remove_tree+0x1eb/0x6d0
    [  168.017346]  remove_tree+0x1eb/0x6d0
    [  168.019969]  remove_tree+0x1eb/0x6d0
    [  168.022128]  qtree_release_dquot+0x291/0x340
    [  168.023297]  v2_release_dquot+0xce/0x120
    [  168.023847]  dquot_release+0x197/0x3e0
    [  168.024358]  ext4_release_dquot+0x22a/0x2d0
    [  168.024932]  dqput.part.0+0x1c9/0x900
    [  168.025430]  __dquot_drop+0x120/0x190
    [  168.025942]  ext4_clear_inode+0x86/0x220
    [  168.026472]  ext4_evict_inode+0x9e8/0xa22
    [  168.028200]  evict+0x29e/0x4f0
    [  168.028625]  dispose_list+0x102/0x1f0
    [  168.029148]  evict_inodes+0x2c1/0x3e0
    [  168.030188]  generic_shutdown_super+0xa4/0x3b0
    [  168.030817]  kill_block_super+0x95/0xd0
    [  168.031360]  deactivate_locked_super+0x85/0xd0
    [  168.031977]  cleanup_mnt+0x2bc/0x480
    [  168.033062]  task_work_run+0xd1/0x170
    [  168.033565]  do_exit+0xa4f/0x2b50
    [  168.037155]  do_group_exit+0xef/0x2d0
    [  168.037666]  __x64_sys_exit_group+0x3a/0x50
    [  168.038237]  do_syscall_64+0x3b/0x90
    [  168.038751]  entry_SYSCALL_64_after_hwframe+0x44/0xae

    In order to reproduce this problem, the following conditions need to be met:
    1. Ext4 filesystem with no journal;
    2. Filesystem image with incorrect quota data;
    3. Abort filesystem forced by user;
    4. umount filesystem;

    As in ext4_quota_write:
    ...
             if (EXT4_SB(sb)->s_journal && !handle) {
                     ext4_msg(sb, KERN_WARNING, "Quota write (off=%llu, len=%llu)"
                             " cancelled because transaction is not started",
                             (unsigned long long)off, (unsigned long long)len);
                     return -EIO;
             }
    ...
    We only check handle if NULL when filesystem has journal. There is need
    check handle if NULL even when filesystem has no journal.

    Signed-off-by: Ye Bin <yebin10@huawei.com>
```

```
    Reviewed-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20211223015506.297766-1-yebin10@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f211541caf77a0941d6c7c34ee5fb78573379366
Author: Luís Henriques <lhenriques@suse.de>
Date:   Tue Dec 14 17:50:58 2021 +0000

    ext4: set csum seed in tmp inode while migrating to extents

    commit e81c9302a6c3c008f5c30beb73b38adb0170ff2d upstream.

    When migrating to extents, the temporary inode will have it's own checksum
    seed.  This means that, when swapping the inodes data, the inode checksums
    will be incorrect.

    This can be fixed by recalculating the extents checksums again.  Or simply
    by copying the seed into the temporary inode.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=213357
    Reported-by: Jeroen van Wolffelaar <jeroen@wolffelaar.nl>
    Signed-off-by: Luís Henriques <lhenriques@suse.de>
    Link: https://lore.kernel.org/r/20211214175058.19511-1-lhenriques@suse.de
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2d5f7996109d3614d90e25cdd7e1fa019f19e3a2
Author: Xin Yin <yinxin.x@bytedance.com>
Date:   Tue Dec 21 10:28:39 2021 +0800

    ext4: fix fast commit may miss tracking range for FALLOC_FL_ZERO_RANGE

    commit 5e4d0eba1ccaf19f93222abdeda5a368be141785 upstream.

    when call falloc with FALLOC_FL_ZERO_RANGE, to set an range to unwritten,
    which has been already initialized. If the range is align to blocksize,
    fast commit will not track range for this change.

    Also track range for unwritten range in ext4_map_blocks().

    Signed-off-by: Xin Yin <yinxin.x@bytedance.com>
    Reviewed-by: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
    Link: https://lore.kernel.org/r/20211221022839.374606-1-yinxin.x@bytedance.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3589522f96e815477ee5f2344f81b91bcad53af2
Author: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
Date:   Wed Dec 1 08:34:21 2021 -0800

    ext4: initialize err_blk before calling __ext4_get_inode_loc

    commit c27c29c6af4f3f4ce925a2111c256733c5a5b430 upstream.

    It is not guaranteed that __ext4_get_inode_loc will definitely set
    err_blk pointer when it returns EIO. To avoid using uninitialized
    variables, let's first set err_blk to 0.

    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Harshad Shirwadkar <harshadshirwadkar@gmail.com>
    Link: https://lore.kernel.org/r/20211201163421.2631661-1-harshads@google.com
```

```
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 56ebfd2b2e784dc76cadc20789eccad2a441239a
Author: Chunguang Xu <brookxu@tencent.com>
Date:   Tue Nov 23 09:17:57 2021 +0800

    ext4: fix a possible ABBA deadlock due to busy PA

    commit 8c80fb312d7abf8bcd66cca1d843a80318a2c522 upstream.

    We found on older kernel (3.10) that in the scenario of insufficient
    disk space, system may trigger an ABBA deadlock problem, it seems that
    this problem still exists in latest kernel, try to fix it here. The
    main process triggered by this problem is that task A occupies the PA
    and waits for the jbd2 transaction finish, the jbd2 transaction waits
    for the completion of task B's IO (plug_list), but task B waits for
    the release of PA by task A to finish discard, which indirectly forms
    an ABBA deadlock. The related calltrace is as follows:

        Task A
        vfs_write
        ext4_mb_new_blocks()
        ext4_mb_mark_diskspace_used()       JBD2
        jbd2_journal_get_write_access()  -> jbd2_journal_commit_transaction()
      ->schedule()                          filemap_fdatawait()
     |                                            |
     | Task B                                     |
     | do_unlinkat()                              |
     | ext4_evict_inode()                         |
     | jbd2_journal_begin_ordered_truncate()      |
     | filemap_fdatawrite_range()                 |
     | ext4_mb_new_blocks()                       |
      -ext4_mb_discard_group_preallocations() <-----

    Here, try to cancel ext4_mb_discard_group_preallocations() internal
    retry due to PA busy, and do a limited number of retries inside
    ext4_mb_discard_preallocations(), which can circumvent the above
    problems, but also has some advantages:

    1. Since the PA is in a busy state, if other groups have free PAs,
       keeping the current PA may help to reduce fragmentation.
    2. Continue to traverse forward instead of waiting for the current
       group PA to be released. In most scenarios, the PA discard time
       can be reduced.

    However, in the case of smaller free space, if only a few groups have
    space, then due to multiple traversals of the group, it may increase
    CPU overhead. But in contrast, I feel that the overall benefit is
    better than the cost.

    Signed-off-by: Chunguang Xu <brookxu@tencent.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Reviewed-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/1637630277-23496-1-git-send-email-brookxu.cn@gmail.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Cc: stable@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 382d022aa664eea134c5103e64329e302ad7a541
Author: Jan Kara <jack@suse.cz>
Date:   Thu Oct 7 17:53:35 2021 +0200

    ext4: make sure quota gets properly shutdown on error
```

```
    commit 15fc69bbbbbc8c72e5f6cc4e1be0f51283c5448e upstream.

    When we hit an error when enabling quotas and setting inode flags, we do
    not properly shutdown quota subsystem despite returning error from
    Q_QUOTAON quotactl. This can lead to some odd situations like kernel
    using quota file while it is still writeable for userspace. Make sure we
    properly cleanup the quota subsystem in case of error.

    Signed-off-by: Jan Kara <jack@suse.cz>
    Cc: stable@kernel.org
    Link: https://lore.kernel.org/r/20211007155336.12493-2-jack@suse.cz
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 232f957560f86d13e81743ab849e4e25766de843
Author: Jan Kara <jack@suse.cz>
Date:   Thu Oct 7 17:53:36 2021 +0200

    ext4: make sure to reset inode lockdep class when quota enabling fails

    commit 4013d47a5307fdb5c13370b5392498b00fedd274 upstream.

    When we succeed in enabling some quota type but fail to enable another
    one with quota feature, we correctly disable all enabled quota types.
    However we forget to reset i_data_sem lockdep class. When the inode gets
    freed and reused, it will inherit this lockdep class (i_data_sem is
    initialized only when a slab is created) and thus eventually lockdep
    barfs about possible deadlocks.

    Reported-and-tested-by: syzbot+3b6f9218b1301ddda3e2@syzkaller.appspotmail.com
    Signed-off-by: Jan Kara <jack@suse.cz>
    Cc: stable@kernel.org
    Link: https://lore.kernel.org/r/20211007155336.12493-3-jack@suse.cz
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4c12d7360d9b28788fa626eabe96adfb620f76c7
Author: Filipe Manana <fdmanana@suse.com>
Date:   Thu Dec 16 15:00:32 2021 +0000

    btrfs: respect the max size in the header when activating swap file

    commit c2f822635df873c510bda6fb7fd1b10b7c31be2d upstream.

    If we extended the size of a swapfile after its header was created (by the
    mkswap utility) and then try to activate it, we will map the entire file
    when activating the swap file, instead of limiting to the max size defined
    in the swap file's header.

    Currently test case generic/643 from fstests fails because we do not
    respect that size limit defined in the swap file's header.

    So fix this by not mapping file ranges beyond the max size defined in the
    swap header.

    This is the same type of bug that iomap used to have, and was fixed in
    commit 36ca7943ac18ae ("mm/swap: consider max pages in
    iomap_swapfile_add_extent").

    Fixes: ed46ff3d423780 ("Btrfs: support swap files")
    CC: stable@vger.kernel.org # 5.4+
    Reviewed-and-tested-by: Josef Bacik <josef@toxicpanda.com
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Reviewed-by: David Sterba <dsterba@suse.com>
```

```
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1aada046526894cd26ad6528560719f495b6d90c
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Wed Dec 8 00:35:49 2021 +0900

    btrfs: zoned: fix chunk allocation condition for zoned allocator

    commit 82187d2ecdfb22ab7ee05f388402a39236d31428 upstream.

    The ZNS specification defines a limit on the number of "active"
    zones. That limit impose us to limit the number of block groups which
    can be used for an allocation at the same time. Not to exceed the
    limit, we reuse the existing active block groups as much as possible
    when we can't activate any other zones without sacrificing an already
    activated block group in commit a85f05e59bc1 ("btrfs: zoned: avoid
    chunk allocation if active block group has enough space").

    However, the check is wrong in two ways. First, it checks the
    condition for every raid index (ffe_ctl->index). Even if it reaches
    the condition and "ffe_ctl->max_extent_size >=
    ffe_ctl->min_alloc_size" is met, there can be other block groups
    having enough space to hold ffe_ctl->num_bytes. (Actually, this won't
    happen in the current zoned code as it only supports SINGLE
    profile. But, it can happen once it enables other RAID types.)

    Second, it checks the active zone availability depending on the
    raid index. The raid index is just an index for
    space_info->block_groups, so it has nothing to do with chunk allocation.

    These mistakes are causing a faulty allocation in a certain
    situation. Consider we are running zoned btrfs on a device whose
    max_active_zone == 0 (no limit). And, suppose no block group have a
    room to fit ffe_ctl->num_bytes but some room to meet
    ffe_ctl->min_alloc_size (i.e. max_extent_size > num_bytes >=
    min_alloc_size).

    In this situation, the following occur:

    - With SINGLE raid_index, it reaches the chunk allocation checking
      code
    - The check returns true because we can activate a new zone (no limit)
    - But, before allocating the chunk, it iterates to the next raid index
      (RAID5)
    - Since there are no RAID5 block groups on zoned mode, it again
      reaches the check code
    - The check returns false because of btrfs_can_activate_zone()'s "if
      (raid_index != BTRFS_RAID_SINGLE)" part
    - That results in returning -ENOSPC without allocating a new chunk

    As a result, we end up hitting -ENOSPC too early.

    Move the check to the right place in the can_allocate_chunk() hook,
    and do the active zone check depending on the allocation flag, not on
    the raid index.

    CC: stable@vger.kernel.org # 5.16
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit df3ebdc229b6a1a0a73d988c73209adc56588574
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Wed Dec 8 00:35:47 2021 +0900
```

```
    btrfs: zoned: unset dedicated block group on allocation failure

    commit 1ada69f61c88abb75a1038ee457633325658a183 upstream.

    Allocating an extent from a block group can fail for various reasons.
    When an allocation from a dedicated block group (for tree-log or
    relocation data) fails, we need to unregister it as a dedicated one so
    that we can allocate a new block group for the dedicated one.

    However, we are returning early when the block group in case it is
    read-only, fully used, or not be able to activate the zone. As a result,
    we keep the non-usable block group as a dedicated one, leading to
    further allocation failure. With many block groups, the allocator will
    iterate hopeless loop to find a free extent, results in a hung task.

    Fix the issue by delaying the return and doing the proper cleanups.

    CC: stable@vger.kernel.org # 5.16
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 00828f2ad2660090d584c35d10f7c034368be83b
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Wed Dec 8 00:35:48 2021 +0900

    btrfs: add extent allocator hook to decide to allocate chunk or not

    commit 50475cd57706359d6cc652be88369dace7a4c2eb upstream.

    Introduce a new hook for an extent allocator policy. With the new
    hook, a policy can decide to allocate a new block group or not. If
    not, it will return -ENOSPC, so btrfs_reserve_extent() will cut the
    allocation size in half and retry the allocation if min_alloc_size is
    large enough.

    The hook has a place holder and will be replaced with the real
    implementation in the next patch.

    CC: stable@vger.kernel.org # 5.16
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f32d02e1a04ec53267f558c0670e4f9b8afbdb40
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Wed Nov 24 14:14:24 2021 -0500

    btrfs: check the root node for uptodate before returning it

    commit 120de408e4b97504a2d9b5ca534b383de2c73d49 upstream.

    Now that we clear the extent buffer uptodate if we fail to write it out
    we need to check to see if our root node is uptodate before we search
    down it.  Otherwise we could return stale data (or potentially corrupt
    data that was caught by the write verification step) and think that the
    path is OK to search down.

    CC: stable@vger.kernel.org # 5.4+
    Reviewed-by: Nikolay Borisov <nborisov@suse.com>
    Signed-off-by: Josef Bacik <josef@toxicpanda.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 0c636ab9078ddebb56c70cbf8a7369e3fa8220df
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Thu Nov 11 14:14:38 2021 +0900

    btrfs: zoned: cache reported zone during mount

    commit 16beac87e95e2fb278b552397c8260637f8a63f7 upstream.

    When mounting a device, we are reporting the zones twice: once for
    checking the zone attributes in btrfs_get_dev_zone_info and once for
    loading block groups' zone info in
    btrfs_load_block_group_zone_info(). With a lot of block groups, that
    leads to a lot of REPORT ZONE commands and slows down the mount
    process.

    This patch introduces a zone info cache in struct
    btrfs_zoned_device_info. The cache is populated while in
    btrfs_get_dev_zone_info() and used for
    btrfs_load_block_group_zone_info() to reduce the number of REPORT ZONE
    commands. The zone cache is then released after loading the block
    groups, as it will not be much effective during the run time.

    Benchmark: Mount an HDD with 57,007 block groups
    Before patch: 171.368 seconds
    After patch: 64.064 seconds

    While it still takes a minute due to the slowness of loading all the
    block groups, the patch reduces the mount time by 1/3.

    Link: https://lore.kernel.org/linux-
btrfs/CAHQ7scUiLtcTqZOMMY5kbWUBOhGRwKo6J6wYPT5WY+C=cD49nQ@mail.gmail.com/
    Fixes: 5b316468983d ("btrfs: get zone information of zoned block devices")
    CC: stable@vger.kernel.org
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b39ece7290e5f5f0e1eb304931644fe47965eca5
Author: Filipe Manana <fdmanana@suse.com>
Date:   Wed Oct 27 18:30:25 2021 +0100

    btrfs: fix deadlock between quota enable and other quota operations

    commit 232796df8c1437c41d308d161007f0715bac0a54 upstream.

    When enabling quotas, we attempt to commit a transaction while holding the
    mutex fs_info->qgroup_ioctl_lock. This can result on a deadlock with other
    quota operations such as:

    - qgroup creation and deletion, ioctl BTRFS_IOC_QGROUP_CREATE;

    - adding and removing qgroup relations, ioctl BTRFS_IOC_QGROUP_ASSIGN.

    This is because these operations join a transaction and after that they
    attempt to lock the mutex fs_info->qgroup_ioctl_lock. Acquiring that mutex
    after joining or starting a transaction is a pattern followed everywhere
    in qgroups, so the quota enablement operation is the one at fault here,
    and should not commit a transaction while holding that mutex.

    Fix this by making the transaction commit while not holding the mutex.
    We are safe from two concurrent tasks trying to enable quotas because
    we are serialized by the rw semaphore fs_info->subvol_sem at
    btrfs_ioctl_quota_ctl(), which is the only call site for enabling
    quotas.
```

```
    When this deadlock happens, it produces a trace like the following:

      INFO: task syz-executor:25604 blocked for more than 143 seconds.
      Not tainted 5.15.0-rc6 #4
      "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
      task:syz-executor state:D stack:24800 pid:25604 ppid: 24873 flags:0x00004004
      Call Trace:
      context_switch kernel/sched/core.c:4940 [inline]
      __schedule+0xcd9/0x2530 kernel/sched/core.c:6287
      schedule+0xd3/0x270 kernel/sched/core.c:6366
      btrfs_commit_transaction+0x994/0x2e90 fs/btrfs/transaction.c:2201
      btrfs_quota_enable+0x95c/0x1790 fs/btrfs/qgroup.c:1120
      btrfs_ioctl_quota_ctl fs/btrfs/ioctl.c:4229 [inline]
      btrfs_ioctl+0x637e/0x7b70 fs/btrfs/ioctl.c:5010
      vfs_ioctl fs/ioctl.c:51 [inline]
      __do_sys_ioctl fs/ioctl.c:874 [inline]
      __se_sys_ioctl fs/ioctl.c:860 [inline]
      __x64_sys_ioctl+0x193/0x200 fs/ioctl.c:860
      do_syscall_x64 arch/x86/entry/common.c:50 [inline]
      do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80
      entry_SYSCALL_64_after_hwframe+0x44/0xae
      RIP: 0033:0x7f86920b2c4d
      RSP: 002b:00007f868f61ac58 EFLAGS: 00000246 ORIG_RAX: 0000000000000010
      RAX: ffffffffffffffda RBX: 00007f86921d90a0 RCX: 00007f86920b2c4d
      RDX: 0000000020005e40 RSI: 00000000c0109428 RDI: 0000000000000008
      RBP: 00007f869212bd80 R08: 0000000000000000 R09: 0000000000000000
      R10: 0000000000000000 R11: 0000000000000246 R12: 00007f86921d90a0
      R13: 00007fff6d233e4f R14: 00007fff6d233ff0 R15: 00007f868f61adc0
      INFO: task syz-executor:25628 blocked for more than 143 seconds.
      Not tainted 5.15.0-rc6 #4
      "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
      task:syz-executor state:D stack:29080 pid:25628 ppid: 24873 flags:0x00004004
      Call Trace:
      context_switch kernel/sched/core.c:4940 [inline]
      __schedule+0xcd9/0x2530 kernel/sched/core.c:6287
      schedule+0xd3/0x270 kernel/sched/core.c:6366
      schedule_preempt_disabled+0xf/0x20 kernel/sched/core.c:6425
      __mutex_lock_common kernel/locking/mutex.c:669 [inline]
      __mutex_lock+0xc96/0x1680 kernel/locking/mutex.c:729
      btrfs_remove_qgroup+0xb7/0x7d0 fs/btrfs/qgroup.c:1548
      btrfs_ioctl_qgroup_create fs/btrfs/ioctl.c:4333 [inline]
      btrfs_ioctl+0x683c/0x7b70 fs/btrfs/ioctl.c:5014
      vfs_ioctl fs/ioctl.c:51 [inline]
      __do_sys_ioctl fs/ioctl.c:874 [inline]
      __se_sys_ioctl fs/ioctl.c:860 [inline]
      __x64_sys_ioctl+0x193/0x200 fs/ioctl.c:860
      do_syscall_x64 arch/x86/entry/common.c:50 [inline]
      do_syscall_64+0x35/0xb0 arch/x86/entry/common.c:80
      entry_SYSCALL_64_after_hwframe+0x44/0xae

    Reported-by: Hao Sun <sunhao.th@gmail.com>
    Link: https://lore.kernel.org/linux-
btrfs/CACkBjsZQF19bQ1C6=yetF3BvL10OSORpFUcWXTP6HErshDB4dQ@mail.gmail.com/
    Fixes: 340f1aa27f36 ("btrfs: qgroups: Move transaction management inside
btrfs_quota_enable/disable")
    CC: stable@vger.kernel.org # 4.19
    Reviewed-by: Qu Wenruo <wqu@suse.com>
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c2f9ea10a1816bd4c54604ecd74aa2f20b636aa0
Author: Nicolas Dichtel <nicolas.dichtel@6wind.com>
Date:   Mon Nov 22 11:33:13 2021 +0100
```

```
    xfrm: fix dflt policy check when there is no policy configured

    commit ec3bb890817e4398f2d46e12e2e205495b116be9 upstream.

    When there is no policy configured on the system, the default policy is
    checked in xfrm_route_forward. However, it was done with the wrong
    direction (XFRM_POLICY_FWD instead of XFRM_POLICY_OUT).
    The default policy for XFRM_POLICY_FWD was checked just before, with a call
    to xfrm[46]_policy_check().

    CC: stable@vger.kernel.org
    Fixes: 2d151d39073a ("xfrm: Add possibility to set the default to block if we have no
policy")
    Signed-off-by: Nicolas Dichtel <nicolas.dichtel@6wind.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b7c44d1329585e7d2e9cd0fae100e154c4fb558c
Author: Ghalem Boudour <ghalem.boudour@6wind.com>
Date:   Fri Nov 19 18:20:16 2021 +0100

    xfrm: fix policy lookup for ipv6 gre packets

    commit bcf141b2eb551b3477b24997ebc09c65f117a803 upstream.

    On egress side, xfrm lookup is called from __gre6_xmit() with the
    fl6_gre_key field not initialized leading to policies selectors check
    failure. Consequently, gre packets are sent without encryption.

    On ingress side, INET6_PROTO_NOPOLICY was set, thus packets were not
    checked against xfrm policies. Like for egress side, fl6_gre_key should be
    correctly set, this is now done in decode_session6().

    Fixes: c12b395a4664 ("gre: Support GRE over IPv6")
    Cc: stable@vger.kernel.org
    Signed-off-by: Ghalem Boudour <ghalem.boudour@6wind.com>
    Signed-off-by: Nicolas Dichtel <nicolas.dichtel@6wind.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 71ceae67ef9b95267d3f1f1d59b7ab9f55f7a709
Author: Pali Rohár <pali@kernel.org>
Date:   Wed Nov 24 16:59:44 2021 +0100

    PCI: pci-bridge-emul: Set PCI_STATUS_CAP_LIST for PCIe device

    commit 3be9d243b21724d49b65043d4520d688b6040b36 upstream.

    Since all PCI Express device Functions are required to implement the PCI
    Express Capability structure, Capabilities List bit in PCI Status Register
    must be hardwired to 1b. Capabilities Pointer register (which is already
    set by pci-bride-emul.c driver) is valid only when Capabilities List is set
    to 1b.

    Link: https://lore.kernel.org/r/20211124155944.1290-7-pali@kernel.org
    Fixes: 23a5fba4d941 ("PCI: Introduce PCI bridge emulated config space common logic")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2c8683fbf14336f7c0332c549bb3b7e87118a5cd
Author: Pali Rohár <pali@kernel.org>
Date:   Wed Nov 24 16:59:43 2021 +0100
```

```
    PCI: pci-bridge-emul: Correctly set PCIe capabilities

    commit 1f1050c5e1fefb34ac90a506b43e9da803b5f8f7 upstream.

    Older mvebu hardware provides PCIe Capability structure only in version 1.
    New mvebu and aardvark hardware provides it in version 2. So do not force
    version to 2 in pci_bridge_emul_init() and rather allow drivers to set
    correct version. Drivers need to set version in pcie_conf.cap field without
    overwriting PCI_CAP_LIST_ID register. Both drivers (mvebu and aardvark) do
    not provide slot support yet, so do not set PCI_EXP_FLAGS_SLOT flag.

    Link: https://lore.kernel.org/r/20211124155944.1290-6-pali@kernel.org
    Fixes: 23a5fba4d941 ("PCI: Introduce PCI bridge emulated config space common logic")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6863f571a54665f71085cb3bba3a8fcfc8214bcd
Author: Pali Rohár <pali@kernel.org>
Date:   Wed Nov 24 16:59:42 2021 +0100

    PCI: pci-bridge-emul: Fix definitions of reserved bits

    commit 12998087d9f48b66965b97412069c7826502cd7e upstream.

    Some bits in PCI_EXP registers are reserved for non-root ports. Driver
    pci-bridge-emul.c implements PCIe Root Port device therefore it should not
    allow setting reserved bits of registers.

    Properly define non-reserved bits for all PCI_EXP registers.

    Link: https://lore.kernel.org/r/20211124155944.1290-5-pali@kernel.org
    Fixes: 23a5fba4d941 ("PCI: Introduce PCI bridge emulated config space common logic")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9e6e6e641f26c5ce23983bd468b9ed82f5b34ea7
Author: Pali Rohár <pali@kernel.org>
Date:   Wed Nov 24 16:59:40 2021 +0100

    PCI: pci-bridge-emul: Properly mark reserved PCIe bits in PCI config space

    commit 7b067ac63a5730d2fae18399fed7e45f23d36912 upstream.

    Some bits in PCI config space are reserved when device is PCIe. Properly
    define behavior of PCI registers for PCIe emulated bridge and ensure that
    it would not be possible change these reserved bits.

    Link: https://lore.kernel.org/r/20211124155944.1290-3-pali@kernel.org
    Fixes: 23a5fba4d941 ("PCI: Introduce PCI bridge emulated config space common logic")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 174a6ab8722eacd08c28581ce20c39bf215331e5
Author: Pali Rohár <pali@kernel.org>
Date:   Wed Nov 24 16:59:39 2021 +0100

    PCI: pci-bridge-emul: Make expansion ROM Base Address register read-only

    commit 1c1a3b4d3e86b997a313ffb297c1129540882859 upstream.
```

```
    If expansion ROM is unsupported (which is the case of pci-bridge-emul.c
    driver) then ROM Base Address register must be implemented as read-only
    register that return 0 when read, same as for unused Base Address
    registers.

    Link: https://lore.kernel.org/r/20211124155944.1290-2-pali@kernel.org
    Fixes: 23a5fba4d941 ("PCI: Introduce PCI bridge emulated config space common logic")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b296d2f6a61a071ae3fdb16dd84cd42fb784b0f9
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Fri Dec 17 15:17:09 2021 +0100

    PCI: pciehp: Use down_read/write_nested(reset_lock) to fix lockdep errors

    commit 085a9f43433f30cbe8a1ade62d9d7827c3217f4d upstream.

    Use down_read_nested() and down_write_nested() when taking the
    ctrl->reset_lock rw-sem, passing the number of PCIe hotplug controllers in
    the path to the PCI root bus as lock subclass parameter.

    This fixes the following false-positive lockdep report when unplugging a
    Lenovo X1C8 from a Lenovo 2nd gen TB3 dock:

      pcieport 0000:06:01.0: pciehp: Slot(1): Link Down
      pcieport 0000:06:01.0: pciehp: Slot(1): Card not present
      ============================================
      WARNING: possible recursive locking detected
      5.16.0-rc2+ #621 Not tainted
      --------------------------------------------
      irq/124-pciehp/86 is trying to acquire lock:
      ffff8e5ac4299ef8 (&ctrl->reset_lock){.+.+}-{3:3}, at: pciehp_check_presence+0x23/0x80

      but task is already holding lock:
      ffff8e5ac4298af8 (&ctrl->reset_lock){.+.+}-{3:3}, at: pciehp_ist+0xf3/0x180

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

             CPU0
             ----
        lock(&ctrl->reset_lock);
        lock(&ctrl->reset_lock);

       *** DEADLOCK ***

       May be due to missing lock nesting notation

      3 locks held by irq/124-pciehp/86:
       #0: ffff8e5ac4298af8 (&ctrl->reset_lock){.+.+}-{3:3}, at: pciehp_ist+0xf3/0x180
       #1: ffffffffa3b024e8 (pci_rescan_remove_lock){+.+.}-{3:3}, at:
pciehp_unconfigure_device+0x31/0x110
       #2: ffff8e5ac1ee2248 (&dev->mutex){....}-{3:3}, at: device_release_driver+0x1c/0x40

      stack backtrace:
      CPU: 4 PID: 86 Comm: irq/124-pciehp Not tainted 5.16.0-rc2+ #621
      Hardware name: LENOVO 20U90SIT19/20U90SIT19, BIOS N2WET30W (1.20 ) 08/26/2021
      Call Trace:
       <TASK>
       dump_stack_lvl+0x59/0x73
       __lock_acquire.cold+0xc5/0x2c6
```

```
        lock_acquire+0xb5/0x2b0
        down_read+0x3e/0x50
        pciehp_check_presence+0x23/0x80
        pciehp_runtime_resume+0x5c/0xa0
        device_for_each_child+0x45/0x70
        pcie_port_device_runtime_resume+0x20/0x30
        pci_pm_runtime_resume+0xa7/0xc0
        __rpm_callback+0x41/0x110
        rpm_callback+0x59/0x70
        rpm_resume+0x512/0x7b0
        __pm_runtime_resume+0x4a/0x90
        __device_release_driver+0x28/0x240
        device_release_driver+0x26/0x40
        pci_stop_bus_device+0x68/0x90
        pci_stop_bus_device+0x2c/0x90
        pci_stop_and_remove_bus_device+0xe/0x20
        pciehp_unconfigure_device+0x6c/0x110
        pciehp_disable_slot+0x5b/0xe0
        pciehp_handle_presence_or_link_change+0xc3/0x2f0
        pciehp_ist+0x179/0x180

    This lockdep warning is triggered because with Thunderbolt, hotplug ports
    are nested. When removing multiple devices in a daisy-chain, each hotplug
    port's reset_lock may be acquired recursively. It's never the same lock, so
    the lockdep splat is a false positive.

    Because locks at the same hierarchy level are never acquired recursively, a
    per-level lockdep class is sufficient to fix the lockdep warning.

    The choice to use one lockdep subclass per pcie-hotplug controller in the
    path to the root-bus was made to conserve class keys because their number
    is limited and the complexity grows quadratically with number of keys
    according to Documentation/locking/lockdep-design.rst.

    Link: https://lore.kernel.org/linux-pci/20190402021933.GA2966@mit.edu/
    Link: https://lore.kernel.org/linux-pci/de684a28-9038-8fc6-27ca-3f6f2f6400d7@redhat.com/
    Link: https://lore.kernel.org/r/20211217141709.379663-1-hdegoede@redhat.com
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=208855
    Reported-by: "Theodore Ts'o" <tytso@mit.edu>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    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 80392a2cb439f31702c8aa79b5994a751843948e
Author: Rob Herring <robh@kernel.org>
Date:   Mon Nov 29 11:36:37 2021 -0600

    PCI: xgene: Fix IB window setup

    commit c7a75d07827a1f33d566e18e6098379cc2a0c2b2 upstream.

    Commit 6dce5aa59e0b ("PCI: xgene: Use inbound resources for setup")
    broke PCI support on XGene. The cause is the IB resources are now sorted
    in address order instead of being in DT dma-ranges order. The result is
    which inbound registers are used for each region are swapped. I don't
    know the details about this h/w, but it appears that IB region 0
    registers can't handle a size greater than 4GB. In any case, limiting
    the size for region 0 is enough to get back to the original assignment
    of dma-ranges to regions.

    Link:
https://lore.kernel.org/all/CA+enf=v9rY_xnZML01oEgKLmvY1NGBUUhnSJaETmXtDtXfaczA@mail.gmail.com/
```

```
    Link: https://lore.kernel.org/r/20211129173637.303201-1-robh@kernel.org
    Fixes: 6dce5aa59e0b ("PCI: xgene: Use inbound resources for setup")
    Reported-by: Stéphane Graber <stgraber@ubuntu.com>
    Tested-by: Stéphane Graber <stgraber@ubuntu.com>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Krzysztof Wilczyński <kw@linux.com>
    Cc: stable@vger.kernel.org # v5.5+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c3f5f3f9d613654d7d65eeb2ddfa534750fd8c6a
Author: José Roberto de Souza <jose.souza@intel.com>
Date:   Thu Jan 13 08:04:37 2022 -0800

    drm/i915/display/ehl: Update voltage swing table

    commit ef3ac01564067a4337bb798b8eddc6ea7b78fd10 upstream.

    EHL table was recently updated with some minor fixes.

    BSpec: 21257
    Cc: stable@vger.kernel.org
    Cc: Clint Taylor <clinton.a.taylor@intel.com>
    Signed-off-by: José Roberto de Souza <jose.souza@intel.com>
    Reviewed-by: Clint Taylor <Clinton.A.Taylor@intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220113160437.49059-1-
jose.souza@intel.com
    (cherry picked from commit 5ec7baef52c367cdbda964aa662f7135c25bab1f)
    Signed-off-by: Tvrtko Ursulin <tvrtko.ursulin@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2b0d7d1abd48e257622c5cea4d6d861596137534
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Fri Jan 7 10:44:17 2022 -0600

    drm/amd/display: Revert W/A for hard hangs on DCN20/DCN21

    commit c4849f88164b13dd141885e28210f599741b304b upstream.

    The WA from commit 2a50edbf10c8 ("drm/amd/display: Apply w/a for hard hang
    on HPD") and commit 1bd3bc745e7f ("drm/amd/display: Extend w/a for hard
    hang on HPD to dcn20") causes a regression in s0ix where the system will
    fail to resume properly on many laptops.  Pull the workarounds out to
    avoid that s0ix regression in the common case.  This HPD hang happens with
    an external device in special circumstances and a new W/A will need to be
    developed for this in the future.

    Cc: stable@vger.kernel.org
    Cc: Qingqing Zhuo <qingqing.zhuo@amd.com>
    Reported-by: Scott Bruce <smbruce@gmail.com>
    Reported-by: Chris Hixon <linux-kernel-bugs@hixontech.com>
    Reported-by: spasswolf@web.de
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=215436
    Link: https://gitlab.freedesktop.org/drm/amd/-/issues/1821
    Link: https://gitlab.freedesktop.org/drm/amd/-/issues/1852
    Fixes: 2a50edbf10c8 ("drm/amd/display: Apply w/a for hard hang on HPD")
    Fixes: 1bd3bc745e7f ("drm/amd/display: Extend w/a for hard hang on HPD to dcn20")
    Reviewed-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d6246a5b649d28c48d9b50733a8331140521bdd1
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Wed Jan 12 22:38:51 2022 -0500
```

```
    drm/amdgpu: don't do resets on APUs which don't support it

    commit e8309d50e97851ff135c4e33325d37b032666b94 upstream.

    It can cause a hang.  This is normally not enabled for GPU
    hangs on these asics, but was recently enabled for handling
    aborted suspends.  This causes hangs on some platforms
    on suspend.

    Fixes: daf8de0874ab5b ("drm/amdgpu: always reset the asic in suspend (v2)")
    Cc: stable@vger.kernel.org
    Bug: https://gitlab.freedesktop.org/drm/amd/-/issues/1858
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c7ba643acdd0f133198604795c94ff95d8552202
Author: Lukas Fink <lukas.fink1@gmail.com>
Date:   Fri Jan 14 07:51:41 2022 +0100

    drm/amdgpu: Fix rejecting Tahiti GPUs

    commit 3993a799fc971bc9b918bd969aa55864447b5dde upstream.

    eb4fd29afd4a ("drm/amdgpu: bind to any 0x1002 PCI diplay class device") added
    generic bindings to amdgpu so that that it binds to all display class devices
    with VID 0x1002 and then rejects those in amdgpu_pci_probe.

    Unfortunately it reuses a driver_data value of 0 to detect those new bindings,
    which is already used to denote CHIP_TAHITI ASICs.

    The driver_data value given to those new bindings was changed in
    dd0761fd24ea1 ("drm/amdgpu: set CHIP_IP_DISCOVERY as the asic type by default")
    to CHIP_IP_DISCOVERY (=36), but it seems that the check in amdgpu_pci_probe
    was forgotten to be changed. Therefore, it still rejects Tahiti GPUs.

    Link: https://gitlab.freedesktop.org/drm/amd/-/issues/1860
    Fixes: eb4fd29afd4a ("drm/amdgpu: bind to any 0x1002 PCI diplay class device")

    Cc: stable@vger.kernel.org
    Signed-off-by: Lukas Fink <lukas.fink1@gmail.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec96690ed667dccd036a713bc9db20370ed231f3
Author: Harry Wentland <harry.wentland@amd.com>
Date:   Tue Jan 4 10:45:41 2022 -0500

    drm/amdgpu: Use correct VIEWPORT_DIMENSION for DCN2

    commit dc5d4aff2e99c312df8abbe1ee9a731d2913bc1b upstream.

    For some reason this file isn't using the appropriate register
    headers for DCN headers, which means that on DCN2 we're getting
    the VIEWPORT_DIMENSION offset wrong.

    This means that we're not correctly carving out the framebuffer
    memory correctly for a framebuffer allocated by EFI and
    therefore see corruption when loading amdgpu before the display
    driver takes over control of the framebuffer scanout.

    Fix this by checking the DCE_HWIP and picking the correct offset
    accordingly.

    Long-term we should expose this info from DC as GMC shouldn't
```

```
    need to know about DCN registers.

    Cc: stable@vger.kernel.org
    Signed-off-by: Harry Wentland <harry.wentland@amd.com>
    Reviewed-by: Huang Rui <ray.huang@amd.com>
    Acked-by: Christian König <christian.koenig@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9998c6b13bb99b87ea8bac4ee5f2e4db9394fbc0
Author: James Smart <jsmart2021@gmail.com>
Date:   Fri Dec 3 16:26:38 2021 -0800

    scsi: lpfc: Fix lpfc_force_rscn ndlp kref imbalance

    commit 7576d48c64f36f6fea9df2882f710a474fa35f40 upstream.

    Issuing lpfc_force_rscn twice results in an ndlp kref use-after-free call
    trace.

    A prior patch reworked the get/put handling by ensuring nlp_get was done
    before WQE submission and a put was done in the completion path.
    Unfortunately, the issue_els_rscn path had a piece of legacy code that did
    a nlp_put, causing an imbalance on the ref counts.

    Fixed by removing the unnecessary legacy code snippet.

    Link: https://lore.kernel.org/r/20211204002644.116455-4-jsmart2021@gmail.com
    Fixes: 4430f7fd09ec ("scsi: lpfc: Rework locations of ndlp reference taking")
    Cc: <stable@vger.kernel.org> # v5.11+
    Co-developed-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: James Smart <jsmart2021@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b9f6c267717caf57b358757bbe164530fdf70680
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Thu Dec 16 20:33:42 2021 +1000

    powerpc/64s/radix: Fix huge vmap false positive

    commit 467ba14e1660b52a2f9338b484704c461bd23019 upstream.

    pmd_huge() is defined to false when HUGETLB_PAGE is not configured, but
    the vmap code still installs huge PMDs. This leads to false bad PMD
    errors when vunmapping because it is not seen as a huge PTE, and the bad
    PMD check catches it. The end result may not be much more serious than
    some bad pmd warning messages, because the pmd_none_or_clear_bad() does
    what we wanted and clears the huge PTE anyway.

    Fix this by checking pmd_is_leaf(), which checks for a PTE regardless of
    config options. The whole huge/large/leaf stuff is a tangled mess but
    that's kernel-wide and not something we can improve much in arch/powerpc
    code.

    pmd_page(), pud_page(), etc., called by vmalloc_to_page() on huge vmaps
    can similarly trigger a false VM_BUG_ON when CONFIG_HUGETLB_PAGE=n, so
    those checks are adjusted. The checks were added by commit d6eacedd1f0e
    ("powerpc/book3s: Use config independent helpers for page table walk"),
    while implementing a similar fix for other page table walking functions.

    Fixes: d909f9109c30 ("powerpc/64s/radix: Enable HAVE_ARCH_HUGE_VMAP")
    Cc: stable@vger.kernel.org # v5.3+
```

```
    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211216103342.609192-1-npiggin@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f5acbab448f82875761eaca10cd401a12405f55d
Author: John David Anglin <dave.anglin@bell.net>
Date:   Wed Dec 22 16:01:31 2021 +0000

    parisc: Fix lpa and lpa_user defines

    commit db19c6f1a2a353cc8dec35b4789733a3cf6e2838 upstream.

    While working on the rewrite to the light-weight syscall and futex code, I
    experimented with using a hash index based on the user physical address of
    atomic variable. This exposed two problems with the lpa and lpa_user defines.

    Because of the copy instruction, the pa argument needs to be an early clobber
    argument. This prevents gcc from allocating the va and pa arguments to the same
    register.

    Secondly, the lpa instruction can cause a page fault so we need to catch
    exceptions.

    Signed-off-by: John David Anglin <dave.anglin@bell.net>
    Fixes: 116d753308cf ("parisc: Use lpa instruction to load physical addresses in driver
code")
    Signed-off-by: Helge Deller <deller@gmx.de>
    Cc: stable@vger.kernel.org # v5.2+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4a36f2f88a6b743421ac9d01952e6344c2077458
Author: Brian Norris <briannorris@chromium.org>
Date:   Wed Nov 3 13:52:00 2021 -0700

    drm/bridge: analogix_dp: Make PSR-exit block less

    commit c4c6ef229593366ab593d4d424addc7025b54a76 upstream.

    Prior to commit 6c836d965bad ("drm/rockchip: Use the helpers for PSR"),
    "PSR exit" used non-blocking analogix_dp_send_psr_spd(). The refactor
    started using the blocking variant, for a variety of reasons -- quoting
    Sean Paul's potentially-faulty memory:

    """
     - To avoid racing a subsequent PSR entry (if exit takes a long time)
     - To avoid racing disable/modeset
     - We're not displaying new content while exiting PSR anyways, so there
       is minimal utility in allowing frames to be submitted
     - We're lying to userspace telling them frames are on the screen when
       we're just dropping them on the floor
    """

    However, I'm finding that this blocking transition is causing upwards of
    60+ ms of unneeded latency on PSR-exit, to the point that initial cursor
    movements when leaving PSR are unbearably jumpy.

    It turns out that we need to meet in the middle somewhere: Sean is right
    that we were "lying to userspace" with a non-blocking PSR-exit, but the
    new blocking behavior is also waiting too long:

    According to the eDP specification, the sink device must support PSR
    entry transitions from both state 4 (ACTIVE_RESYNC) and state 0
    (INACTIVE). It also states that in ACTIVE_RESYNC, "the Sink device must
    display the incoming active frames from the Source device with no
```

```
    visible glitches and/or artifacts."

    Thus, for our purposes, we only need to wait for ACTIVE_RESYNC before
    moving on; we are ready to display video, and subsequent PSR-entry is
    safe.

    Tested on a Samsung Chromebook Plus (i.e., Rockchip RK3399 Gru Kevin),
    where this saves about 60ms of latency, for PSR-exit that used to
    take about 80ms.

    Fixes: 6c836d965bad ("drm/rockchip: Use the helpers for PSR")
    Cc: <stable@vger.kernel.org>
    Cc: Zain Wang <wzz@rock-chips.com>
    Cc: Tomasz Figa <tfiga@chromium.org>
    Cc: Heiko Stuebner <heiko@sntech.de>
    Cc: Sean Paul <seanpaul@chromium.org>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reviewed-by: Sean Paul <seanpaul@chromium.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20211103135112.v3.1.I67612ea073c3306c71b46a87b
e894f79707082df@changeid
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fc10fbe2573e8f66d3700e684fc0b3de06a8b1d1
Author: Ilia Mirkin <imirkin@alum.mit.edu>
Date:   Sun Mar 7 12:48:53 2021 -0500

    drm/nouveau/kms/nv04: use vzalloc for nv04_display

    commit bd6e07e72f37f34535bec7eebc807e5fcfe37b43 upstream.

    The struct is giant, and triggers an order-7 allocation (512K). There is
    no reason for this to be kmalloc-type memory, so switch to vmalloc. This
    should help loading nouveau on low-memory and/or long-running systems.

    Reported-by: Nathan E. Egge <unlord@xiph.org>
    Signed-off-by: Ilia Mirkin <imirkin@alum.mit.edu>
    Cc: stable@vger.kernel.org
    Signed-off-by: Ben Skeggs <bskeggs@redhat.com>
    Reviewed-by: Karol Herbst <kherbst@redhat.com>
    Signed-off-by: Karol Herbst <kherbst@redhat.com>
    Link: https://gitlab.freedesktop.org/drm/nouveau/-/merge_requests/10
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 950eddff40bfe1db759e3ddc41f03266ca9bf135
Author: Yizhuo Zhai <yzhai003@ucr.edu>
Date:   Fri Dec 17 20:22:23 2021 -0800

    drm/amd/display: Fix the uninitialized variable in enable_stream_features()

    commit 0726ed3065eeb910f9cea0c933bc021a848e00b3 upstream.

    In function enable_stream_features(), the variable "old_downspread.raw"
    could be uninitialized if core_link_read_dpcd() fails, however, it is
    used in the later if statement, and further, core_link_write_dpcd()
    may write random value, which is potentially unsafe.

    Fixes: 6016cd9dba0f ("drm/amd/display: add helper for enabling mst stream features")
    Cc: stable@vger.kernel.org
    Signed-off-by: Yizhuo Zhai <yzhai003@ucr.edu>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit deb03f4c3c8aa350e332ba19d4e40faac36ce646
```

```
Author: Lucas Stach <l.stach@pengutronix.de>
Date:   Fri Dec 17 11:59:28 2021 +0100

    drm/etnaviv: limit submit sizes

    commit 6dfa2fab8ddd46faa771a102672176bee7a065de upstream.

    Currently we allow rediculous amounts of kernel memory being allocated
    via the etnaviv GEM_SUBMIT ioctl, which is a pretty easy DoS vector. Put
    some reasonable limits in to fix this.

    The commandstream size is limited to 64KB, which was already a soft limit
    on older kernels after which the kernel only took submits on a best effort
    base, so there is no userspace that tries to submit commandstreams larger
    than this. Even if the whole commandstream is a single incrementing address
    load, the size limit also limits the number of potential relocs and
    referenced buffers to slightly under 64K, so use the same limit for those
    arguments. The performance monitoring infrastructure currently supports
    less than 50 performance counter signals, so limiting them to 128 on a
    single submit seems like a reasonably future-proof number for now. This
    number can be bumped if needed without breaking the interface.

    Cc: stable@vger.kernel.org
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Lucas Stach <l.stach@pengutronix.de>
    Reviewed-by: Christian Gmeiner <christian.gmeiner@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb6f9bfac1af3d7521beb6834839f988776c9d87
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Wed Dec 1 02:23:17 2021 +0300

    drm/tegra: submit: Add missing pm_runtime_mark_last_busy()

    commit a21115dd38c6cf396ba39aefd561e7903ca6149d upstream.

    Runtime PM auto-suspension doesn't work without pm_runtime_mark_last_busy(),
    add it.

    Cc: <stable@vger.kernel.org>
    Reviewed-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ab9559f8ff5019e6af77922354276c6471e03be3
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Wed Dec 1 14:59:29 2021 +0200

    device property: Fix fwnode_graph_devcon_match() fwnode leak

    commit 4a7f4110f79163fd53ea65438041994ed615e3af upstream.

    For each endpoint it encounters, fwnode_graph_devcon_match() checks
    whether the endpoint's remote port parent device is available. If it is
    not, it ignores the endpoint but does not put the reference to the remote
    endpoint port parent fwnode. For available devices the fwnode handle
    reference is put as expected.

    Put the reference for unavailable devices now.

    Fixes: 637e9e52b185 ("device connection: Find device connections also from device
graphs")
    Cc: 5.1+ <stable@vger.kernel.org> # 5.1+
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
```

```
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b41054c73268eb75a8e4dac85dd425b2ca910fd4
Author: Alexander Gordeev <agordeev@linux.ibm.com>
Date:   Thu Nov 4 07:14:44 2021 +0100

    s390/mm: fix 2KB pgtable release race

    commit c2c224932fd0ee6854d6ebfc8d059c2bcad86606 upstream.

    There is a race on concurrent 2KB-pgtables release paths when
    both upper and lower halves of the containing parent page are
    freed, one via page_table_free_rcu() + __tlb_remove_table(),
    and the other via page_table_free(). The race might lead to a
    corruption as result of remove of list item in page_table_free()
    concurrently with __free_page() in __tlb_remove_table().

    Let's assume first the lower and next the upper 2KB-pgtables are
    freed from a page. Since both halves of the page are allocated
    the tracking byte (bits 24-31 of the page _refcount) has value
    of 0x03 initially:

    CPU0                            CPU1
    ----                            ----

    page_table_free_rcu() // lower half
    {
            // _refcount[31..24] == 0x03
            ...
            atomic_xor_bits(&page->_refcount,
                            0x11U << (0 + 24));
            // _refcount[31..24] <= 0x12
            ...
            table = table | (1U << 0);
            tlb_remove_table(tlb, table);
    }
    ...
    __tlb_remove_table()
    {
            // _refcount[31..24] == 0x12
            mask = _table & 3;
            // mask <= 0x01
            ...

                                    page_table_free() // upper half
                                    {
                                            // _refcount[31..24] == 0x12
                                            ...
                                            atomic_xor_bits(
                                                    &page->_refcount,
                                                    1U << (1 + 24));
                                            // _refcount[31..24] <= 0x10
                                            // mask <= 0x10
                                            ...
            atomic_xor_bits(&page->_refcount,
                            mask << (4 + 24));
            // _refcount[31..24] <= 0x00
            // mask <= 0x00
            ...
            if (mask != 0) // == false
                    break;
            fallthrough;
            ...
```

```
                                        if (mask & 3) // == false
                                                ...
                                        else
        __free_page(page);                      list_del(&page->lru);
        ^^^^^^^^^^^^^^^^^^      RACE!           ^^^^^^^^^^^^^^^^^^^^^
    }                                   ...
                                    }
```

    The problem is page_table_free() releases the page as result of
    lower nibble unset and __tlb_remove_table() observing zero too
    early. With this update page_table_free() will use the similar
    logic as page_table_free_rcu() + __tlb_remove_table(), and mark
    the fragment as pending for removal in the upper nibble until
    after the list_del().

    In other words, the parent page is considered as unreferenced and
    safe to release only when the lower nibble is cleared already and
    unsetting a bit in upper nibble results in that nibble turned zero.

    Cc: stable@vger.kernel.org
    Suggested-by: Vlastimil Babka <vbabka@suse.com>
    Reviewed-by: Gerald Schaefer <gerald.schaefer@linux.ibm.com>
    Signed-off-by: Alexander Gordeev <agordeev@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cff07a2d07a41bac71b6d29bc638b22c9b6e8211
Author: Ilan Peer <ilan.peer@intel.com>
Date: Fri Dec 10 09:06:21 2021 +0200

    iwlwifi: mvm: Increase the scan timeout guard to 30 seconds

    commit ced50f1133af12f7521bb777fcf4046ca908fb77 upstream.

    With the introduction of 6GHz channels the scan guard timeout should
    be adjusted to account for the following extreme case:

    - All 6GHz channels are scanned passively: 58 channels.
    - The scan is fragmented with the following parameters: 3 fragments,
      95 TUs suspend time, 44 TUs maximal out of channel time.

    The above would result with scan time of more than 24 seconds. Thus,
    set the timeout to 30 seconds.

    Cc: stable@vger.kernel.org
    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.20211210090244.3c851b93aef5.I346fa2e1d79220a6770496e773c6f
87a2ad9e6c4@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8aed51bd73fad9b01ec5f57a65e7607b6d795ebd
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date: Sat Oct 16 08:44:28 2021 +0200

    remoteproc: imx_rproc: Fix a resource leak in the remove function

    commit 4da96175014be67c846fd274eace08066e525d75 upstream.

    'priv->workqueue' is destroyed in the error handling path of the probe but
    not in the remove function.

    Add the missing call to release some resources.

```
    Cc: stable <stable@vger.kernel.org>
    Fixes: 2df7062002d0 ("remoteproc: imx_proc: enable virtio/mailbox")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Reviewed-by: Peng Fan <peng.fan@nxp.com>
    Tested-by: Peng Fan <peng.fan@nxp.com>
    Link:
https://lore.kernel.org/r/d28ca94a4031bd7297d47c2164e18885a5a6ec19.1634366546.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f79e9f32eee8f570ad01ba104f3cbcea222e7b38
Author: Steven Rostedt <rostedt@goodmis.org>
Date:   Fri Jan 7 17:56:56 2022 -0500

    tracing: Have syscall trace events use trace_event_buffer_lock_reserve()

    commit 3e2a56e6f639492311e0a8533f0a7aed60816308 upstream.

    Currently, the syscall trace events call trace_buffer_lock_reserve()
    directly, which means that it misses out on some of the filtering
    optimizations provided by the helper function
    trace_event_buffer_lock_reserve(). Have the syscall trace events call that
    instead, as it was missed when adding the update to use the temp buffer
    when filtering.

    Link: https://lkml.kernel.org/r/20220107225839.823118570@goodmis.org

    Cc: stable@vger.kernel.org
    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: Andrew Morton <akpm@linux-foundation.org>
    Cc: Tom Zanussi <zanussi@kernel.org>
    Reviewed-by: Masami Hiramatsu <mhiramat@kernel.org>
    Fixes: 0fc1b09ff1ff4 ("tracing: Use temp buffer when filtering events")
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7cf6dacdd0b14490904c948fd9bb111e498a53c1
Author: Xiangyang Zhang <xyz.sun.ok@gmail.com>
Date:   Fri Jan 7 23:02:42 2022 +0800

    tracing/kprobes: 'nmissed' not showed correctly for kretprobe

    commit dfea08a2116fe327f79d8f4d4b2cf6e0c88be11f upstream.

    The 'nmissed' column of the 'kprobe_profile' file for kretprobe is
    not showed correctly, kretprobe can be skipped by two reasons,
    shortage of kretprobe_instance which is counted by tk->rp.nmissed,
    and kprobe itself is missed by some reason, so to show the sum.

    Link: https://lkml.kernel.org/r/20220107150242.5019-1-xyz.sun.ok@gmail.com

    Cc: stable@vger.kernel.org
    Fixes: 4a846b443b4e ("tracing/kprobes: Cleanup kprobe tracer code")
    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Signed-off-by: Xiangyang Zhang <xyz.sun.ok@gmail.com>
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6de8ecd68f32d7acf5eb8a6f64a76607fa6dfb91
Author: Nikita Yushchenko <nikita.yushchenko@virtuozzo.com>
Date:   Sun Jan 9 18:34:59 2022 +0300

    tracing/osnoise: Properly unhook events if start_per_cpu_kthreads() fails
```

```
    commit 0878355b51f5f26632e652c848a8e174bb02d22d upstream.

    If start_per_cpu_kthreads() called from osnoise_workload_start() returns
    error, event hooks are left in broken state: unhook_irq_events() called
    but unhook_thread_events() and unhook_softirq_events() not called, and
    trace_osnoise_callback_enabled flag not cleared.

    On the next tracer enable, hooks get not installed due to
    trace_osnoise_callback_enabled flag.

    And on the further tracer disable an attempt to remove non-installed
    hooks happened, hitting a WARN_ON_ONCE() in tracepoint_remove_func().

    Fix the error path by adding the missing part of cleanup.
    While at this, introduce osnoise_unhook_events() to avoid code
    duplication between this error path and normal tracer disable.

    Link: https://lkml.kernel.org/r/20220109153459.3701773-1-nikita.yushchenko@virtuozzo.com

    Cc: stable@vger.kernel.org
    Fixes: bce29ac9ce0b ("trace: Add osnoise tracer")
    Acked-by: Daniel Bristot de Oliveira <bristot@kernel.org>
    Signed-off-by: Nikita Yushchenko <nikita.yushchenko@virtuozzo.com>
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7dd8db89e42d50fdb241de89509b2fb00a7d89d4
Author: Andrey Ryabinin <arbn@yandex-team.com>
Date:   Mon Nov 15 19:46:06 2021 +0300

    sched/cpuacct: Fix user/system in shown cpuacct.usage*

    commit dd02d4234c9a2214a81c57a16484304a1a51872a upstream.

    cpuacct has 2 different ways of accounting and showing user
    and system times.

    The first one uses cpuacct_account_field() to account times
    and cpuacct.stat file to expose them. And this one seems to work ok.

    The second one is uses cpuacct_charge() function for accounting and
    set of cpuacct.usage* files to show times. Despite some attempts to
    fix it in the past it still doesn't work. Sometimes while running KVM
    guest the cpuacct_charge() accounts most of the guest time as
    system time. This doesn't match with user&system times shown in
    cpuacct.stat or proc/<pid>/stat.

    Demonstration:
     # git clone https://github.com/aryabinin/kvmsample
     # make
     # mkdir /sys/fs/cgroup/cpuacct/test
     # echo $$ > /sys/fs/cgroup/cpuacct/test/tasks
     # ./kvmsample &
     # for i in {1..5}; do cat /sys/fs/cgroup/cpuacct/test/cpuacct.usage_sys; sleep 1; done
     1976535645
     2979839428
     3979832704
     4983603153
     5983604157

    Use cpustats accounted in cpuacct_account_field() as the source
    of user/sys times for cpuacct.usage* files. Make cpuacct_charge()
    to account only summary execution time.
```

```
    Fixes: d740037fac70 ("sched/cpuacct: Split usage accounting into user_usage and
sys_usage")
    Signed-off-by: Andrey Ryabinin <arbn@yandex-team.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Daniel Jordan <daniel.m.jordan@oracle.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20211115164607.23784-3-arbn@yandex-team.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 92b0ef7d669b1bdde4af045f608137e986fa2c68
Author: Andrey Ryabinin <arbn@yandex-team.com>
Date:   Mon Nov 15 19:46:04 2021 +0300

    cputime, cpuacct: Include guest time in user time in cpuacct.stat

    commit 9731698ecb9c851f353ce2496292ff9fcea39dff upstream.

    cpuacct.stat in no-root cgroups shows user time without guest time
    included int it. This doesn't match with user time shown in root
    cpuacct.stat and /proc/<pid>/stat. This also affects cgroup2's cpu.stat
    in the same way.

    Make account_guest_time() to add user time to cgroup's cpustat to
    fix this.

    Fixes: ef12fefabf94 ("cpuacct: add per-cgroup utime/stime statistics")
    Signed-off-by: Andrey Ryabinin <arbn@yandex-team.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Daniel Jordan <daniel.m.jordan@oracle.com>
    Acked-by: Tejun Heo <tj@kernel.org>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20211115164607.23784-1-arbn@yandex-team.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 579718f1aa15ef010a27aab9d07f7fec342ab50a
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sat Dec 18 10:58:56 2021 +0100

    serial: Fix incorrect rs485 polarity on uart open

    commit d3b3404df318504ec084213ab1065b73f49b0f1d upstream.

    Commit a6845e1e1b78 ("serial: core: Consider rs485 settings to drive
    RTS") sought to deassert RTS when opening an rs485-enabled uart port.
    That way, the transceiver does not occupy the bus until it transmits
    data.

    Unfortunately, the commit mixed up the logic and *asserted* RTS instead
    of *deasserting* it:

    The commit amended uart_port_dtr_rts(), which raises DTR and RTS when
    opening an rs232 port.  "Raising" actually means lowering the signal
    that's coming out of the uart, because an rs232 transceiver not only
    changes a signal's voltage level, it also *inverts* the signal.  See
    the simplified schematic in the MAX232 datasheet for an example:
    https://www.ti.com/lit/ds/symlink/max232.pdf

    So, to raise RTS on an rs232 port, TIOCM_RTS is *set* in port->mctrl
    and that results in the signal being driven low.

    In contrast to rs232, the signal level for rs485 Transmit Enable is the
    identity, not the inversion:  If the transceiver expects a "high" RTS
    signal for Transmit Enable, the signal coming out of the uart must also
    be high, so TIOCM_RTS must be *cleared* in port->mctrl.
```

```
    The commit did the exact opposite, but it's easy to see why given the
    confusing semantics of rs232 and rs485.  Fix it.

    Fixes: a6845e1e1b78 ("serial: core: Consider rs485 settings to drive RTS")
    Cc: stable@vger.kernel.org # v4.14+
    Cc: Rafael Gago Castano <rgc@hms.se>
    Cc: Jan Kiszka <jan.kiszka@siemens.com>
    Cc: Su Bao Cheng <baocheng.su@siemens.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Link:
https://lore.kernel.org/r/9395767847833f2f3193c49cde38501eeb3b5669.1639821059.git.lukas@wunn
er.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a5546d1bb58f03594c602852cf4e7482f9d8c22e
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Mon Nov 22 17:05:31 2021 +0800

    fuse: Pass correct lend value to filemap_write_and_wait_range()

    commit e388164ea385f04666c4633f5dc4f951fca71890 upstream.

    The acceptable maximum value of lend parameter in
    filemap_write_and_wait_range() is LLONG_MAX rather than -1. And there is
    also some logic depending on LLONG_MAX check in write_cache_pages(). So
    let's pass LLONG_MAX to filemap_write_and_wait_range() in
    fuse_writeback_range() instead.

    Fixes: 59bda8ecee2f ("fuse: flush extending writes")
    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Cc: <stable@vger.kernel.org> # v5.15
    Signed-off-by: Miklos Szeredi <mszeredi@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4266bbb97e150126ecc8b73f07f80b4bf9fb7be3
Author: Borislav Petkov <bp@suse.de>
Date:   Fri Dec 17 16:49:25 2021 +0100

    x86/mce: Check regs before accessing it

    commit 1acd85feba81084fcef00b73fc1601e42b77c5d8 upstream.

    Commit in Fixes accesses pt_regs before checking whether it is NULL or
    not. Make sure the NULL pointer check happens first.

    Fixes: 0a5b288e85bb ("x86/mce: Prevent severity computation from being instrumented")
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Tony Luck <tony.luck@intel.com>
    Link: https://lore.kernel.org/r/20211217102029.GA29708@kili
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b439dda3160bb00737e61beec1493cdbf3e63d2e
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Tue Dec 28 22:09:17 2021 +0100

    HID: magicmouse: Fix an error handling path in magicmouse_probe()

    commit 33812fc7c8d77a43b7e2bf36a0d5a57c277a4b0c upstream.

    If the timer introduced by the commit below is started, then it must be
    deleted in the error handling of the probe. Otherwise it would trigger
    once the driver is no more.
```

```
    Fixes: 0b91b4e4dae6 ("HID: magicmouse: Report battery level over USB")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Tested-by: José Expósito <jose.exposito89@gmail.com>
    Reported-by: <syzbot+a437546ec71b04dfb5ac@syzkaller.appspotmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 92053ae0398a58c6914e807d5e070297ba8cfe88
Author: Xiao Ni <xni@redhat.com>
Date:   Fri Dec 10 17:31:15 2021 +0800

    md: Move alloc/free acct bioset in to personality

    commit 0c031fd37f69deb0cd8c43bbfcfccd62ebd7e952 upstream.

    bioset acct is only needed for raid0 and raid5. Therefore, md_run only
    allocates it for raid0 and raid5. However, this does not cover
    personality takeover, which may cause uninitialized bioset. For example,
    the following repro steps:

      mdadm -CR /dev/md0 -l1 -n2 /dev/loop0 /dev/loop1
      mdadm --wait /dev/md0
      mkfs.xfs /dev/md0
      mdadm /dev/md0 --grow -l5
      mount /dev/md0 /mnt

    causes panic like:

    [  225.933939] BUG: kernel NULL pointer dereference, address: 0000000000000000
    [  225.934903] #PF: supervisor instruction fetch in kernel mode
    [  225.935639] #PF: error_code(0x0010) - not-present page
    [  225.936361] PGD 0 P4D 0
    [  225.936677] Oops: 0010 [#1] PREEMPT SMP DEBUG_PAGEALLOC KASAN PTI
    [  225.937525] CPU: 27 PID: 1133 Comm: mount Not tainted 5.16.0-rc3+ #706
    [  225.938416] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.13.0-
2.module_el8.4.0+547+a85d02ba 04/01/2014
    [  225.939922] RIP: 0010:0x0
    [  225.940289] Code: Unable to access opcode bytes at RIP 0xffffffffffffffd6.
    [  225.941196] RSP: 0018:ffff88815897eff0 EFLAGS: 00010246
    [  225.941897] RAX: 0000000000000000 RBX: 0000000000092800 RCX: ffffffff81370a39
    [  225.942813] RDX: dfffc0000000000 RSI: 0000000000000000 RDI: 0000000000092800
    [  225.943772] RBP: 1ffff1102b12fe04 R08: fffffbfff0b43c01 R09: fffffbfff0b43c01
    [  225.944807] R10: ffffffff85a1e007 R11: fffffbfff0b43c00 R12: ffff88810eaaaf58
    [  225.945757] R13: 0000000000000000 R14: ffff88810eaaafb8 R15: ffff88815897f040
    [  225.946709] FS:  00007ff3f2505080(0000) GS:ffff888fb5e00000(0000)
knlGS:0000000000000000
    [  225.947814] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [  225.948556] CR2: ffffffffffffffd6 CR3: 000000015aa5a006 CR4: 0000000000370ee0
    [  225.949537] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [  225.950455] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [  225.951414] Call Trace:
    [  225.951787]  <TASK>
    [  225.952120]  mempool_alloc+0xe5/0x250
    [  225.952625]  ? mempool_resize+0x370/0x370
    [  225.953187]  ? rcu_read_lock_sched_held+0xa1/0xd0
    [  225.953862]  ? rcu_read_lock_bh_held+0xb0/0xb0
    [  225.954464]  ? sched_clock_cpu+0x15/0x120
    [  225.955019]  ? find_held_lock+0xac/0xd0
    [  225.955564]  bio_alloc_bioset+0x1ed/0x2a0
    [  225.956080]  ? lock_downgrade+0x3a0/0x3a0
    [  225.956644]  ? bvec_alloc+0xc0/0xc0
    [  225.957135]  bio_clone_fast+0x19/0x80
    [  225.957651]  raid5_make_request+0x1370/0x1b70
    [  225.958286]  ? sched_clock_cpu+0x15/0x120
    [  225.958797]  ? __lock_acquire+0x8b2/0x3510
```

```
[  225.959339]  ? raid5_get_active_stripe+0xce0/0xce0
[  225.959986]  ? lock_is_held_type+0xd8/0x130
[  225.960528]  ? rcu_read_lock_sched_held+0xa1/0xd0
[  225.961135]  ? rcu_read_lock_bh_held+0xb0/0xb0
[  225.961703]  ? sched_clock_cpu+0x15/0x120
[  225.962232]  ? lock_release+0x27a/0x6c0
[  225.962746]  ? do_wait_intr_irq+0x130/0x130
[  225.963302]  ? lock_downgrade+0x3a0/0x3a0
[  225.963815]  ? lock_release+0x6c0/0x6c0
[  225.964348]  md_handle_request+0x342/0x530
[  225.964888]  ? set_in_sync+0x170/0x170
[  225.965397]  ? blk_queue_split+0x133/0x150
[  225.965988]  ? __blk_queue_split+0x8b0/0x8b0
[  225.966524]  ? submit_bio_checks+0x3b2/0x9d0
[  225.967069]  md_submit_bio+0x127/0x1c0
[...]

Fix this by moving alloc/free of acct bioset to pers->run and pers->free.

While we are on this, properly handle md_integrity_register() error in
raid0_run().

Fixes: daee2024715d (md: check level before create and exit io_acct_set)
Cc: stable@vger.kernel.org
Acked-by: Guoqing Jiang <guoqing.jiang@linux.dev>
Signed-off-by: Xiao Ni <xni@redhat.com>
Signed-off-by: Song Liu <song@kernel.org>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 562977d039f2d9f2ee225d073a6f496bf307325d
Author: Oleksandr Andrushchenko <oleksandr_andrushchenko@epam.com>
Date:   Fri Dec 10 11:28:17 2021 +0200

    xen/gntdev: fix unmap notification order

    commit ce2f46f3531a03781181b7f4bd1ff9f8c5086e7e upstream.

    While working with Xen's libxenvchan library I have faced an issue with
    unmap notifications sent in wrong order if both UNMAP_NOTIFY_SEND_EVENT
    and UNMAP_NOTIFY_CLEAR_BYTE were requested: first we send an event channel
    notification and then clear the notification byte which renders in the below
    inconsistency (cli_live is the byte which was requested to be cleared on unmap):

    [  444.514243] gntdev_put_map UNMAP_NOTIFY_SEND_EVENT map->notify.event 6
    libxenvchan_is_open cli_live 1
    [  444.515239] __unmap_grant_pages UNMAP_NOTIFY_CLEAR_BYTE at 14

    Thus it is not possible to reliably implement the checks like
    - wait for the notification (UNMAP_NOTIFY_SEND_EVENT)
    - check the variable (UNMAP_NOTIFY_CLEAR_BYTE)
    because it is possible that the variable gets checked before it is cleared
    by the kernel.

    To fix that we need to re-order the notifications, so the variable is first
    gets cleared and then the event channel notification is sent.
    With this fix I can see the correct order of execution:

    [   54.522611] __unmap_grant_pages UNMAP_NOTIFY_CLEAR_BYTE at 14
    [   54.537966] gntdev_put_map UNMAP_NOTIFY_SEND_EVENT map->notify.event 6
    libxenvchan_is_open cli_live 0

    Cc: stable@vger.kernel.org
    Signed-off-by: Oleksandr Andrushchenko <oleksandr_andrushchenko@epam.com>
    Reviewed-by: Boris Ostrovsky <boris.ostrovsky@oracle.com>
    Link: https://lore.kernel.org/r/20211210092817.580718-1-andr2000@gmail.com
```

```
    Signed-off-by: Juergen Gross <jgross@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 71131329be7826361c10ac40809ed9333689a9f7
Author: Kunihiko Hayashi <hayashi.kunihiko@socionext.com>
Date:   Wed Dec 22 13:48:12 2021 +0900

    spi: uniphier: Fix a bug that doesn't point to private data correctly

    commit 80bb73a9fbcde4ecc55e12f10c73fabbe68a24d1 upstream.

    In uniphier_spi_remove(), there is a wrong code to get private data from
    the platform device, so the driver can't be removed properly.

    The driver should get spi_master from the platform device and retrieve
    the private data from it.

    Cc: <stable@vger.kernel.org>
    Fixes: 5ba155a4d4cc ("spi: add SPI controller driver for UniPhier SoC")
    Signed-off-by: Kunihiko Hayashi <hayashi.kunihiko@socionext.com>
    Link: https://lore.kernel.org/r/1640148492-32178-1-git-send-email-
hayashi.kunihiko@socionext.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7c02e1360b2d6d74973a0018bdfe7392af320b3b
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Wed Nov 24 22:01:04 2021 +0300

    mfd: tps65910: Set PWR_OFF bit during driver probe

    commit 7620ad0bdfac1efff4a1228cd36ae62a9d8206b0 upstream.

    The PWR_OFF bit needs to be set in order to power off properly, without
    hanging PMIC. This bit needs to be set early in order to allow thermal
    protection of NVIDIA Terga SoCs to power off hardware properly, otherwise
    a battery re-plug may be needed on some devices to recover after the hang.

    Cc: <stable@vger.kernel.org>
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Tested-by: Svyatoslav Ryhel <clamor95@gmail.com> # ASUS TF201
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211124190104.23554-1-digetx@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 13fe4bdb76ca9ec1e441e6292220dbd55fef07cc
Author: Patrick Williams <patrick@stwcx.xyz>
Date:   Thu Dec 23 09:49:31 2021 -0600

    tpm: fix NPE on probe for missing device

    commit 84cc69589700b90a4c8d27b481a51fce8cca6051 upstream.

    When using the tpm_tis-spi driver on a system missing the physical TPM,
    a null pointer exception was observed.

        [    0.938677] Unable to handle kernel NULL pointer dereference at virtual address
00000004
        [    0.939020] pgd = 10c753cb
        [    0.939237] [00000004] *pgd=00000000
        [    0.939808] Internal error: Oops: 5 [#1] SMP ARM
        [    0.940157] CPU: 0 PID: 48 Comm: kworker/u4:1 Not tainted 5.15.10-dd1e40c #1
        [    0.940364] Hardware name: Generic DT based system
        [    0.940601] Workqueue: events_unbound async_run_entry_fn
        [    0.941048] PC is at tpm_tis_remove+0x28/0xb4
```

```
         [    0.941196] LR is at tpm_tis_core_init+0x170/0x6ac

    This is due to an attempt in 'tpm_tis_remove' to use the drvdata, which
    was not initialized in 'tpm_tis_core_init' prior to the first error.

    Move the initialization of drvdata earlier so 'tpm_tis_remove' has
    access to it.

    Signed-off-by: Patrick Williams <patrick@stwcx.xyz>
    Fixes: 79ca6f74dae0 ("tpm: fix Atmel TPM crash caused by too frequent queries")
    Cc: stable@vger.kernel.org
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b66371ef99a064412478909098576fbc0b37dc4a
Author: Lino Sanfilippo <LinoSanfilippo@gmx.de>
Date:   Mon Dec 20 16:06:35 2021 +0100

    tpm: fix potential NULL pointer access in tpm_del_char_device

    commit eabad7ba2c752392ae50f24a795093fb115b686d upstream.

    Some SPI controller drivers unregister the controller in the shutdown
    handler (e.g. BCM2835). If such a controller is used with a TPM 2 slave
    chip->ops may be accessed when it is already NULL:

    At system shutdown the pre-shutdown handler tpm_class_shutdown() shuts down
    TPM 2 and sets chip->ops to NULL. Then at SPI controller unregistration
    tpm_tis_spi_remove() is called and eventually calls tpm_del_char_device()
    which tries to shut down TPM 2 again. Thereby it accesses chip->ops again:
    (tpm_del_char_device calls tpm_chip_start which calls tpm_clk_enable which
    calls chip->ops->clk_enable).

    Avoid the NULL pointer access by testing if chip->ops is valid and skipping
    the TPM 2 shutdown procedure in case it is NULL.

    Cc: stable@vger.kernel.org
    Signed-off-by: Lino Sanfilippo <LinoSanfilippo@gmx.de>
    Fixes: 39d0099f9439 ("powerpc/pseries: Add shutdown() to vio_driver and vio_bus")
    Reviewed-by: Stefan Berger <stefanb@linux.ibm.com>
    Tested-by: Stefan Berger <stefanb@linux.ibm.com>
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9ce5afd05eb178e0d79f4153022247d8f52e58e6
Author: Petr Cvachoucek <cvachoucek@gmail.com>
Date:   Mon Aug 30 21:20:37 2021 +0200

    ubifs: Error path in ubifs_remount_rw() seems to wrongly free write buffers

    commit 3fea4d9d160186617ff40490ae01f4f4f36b28ff upstream.

    it seems freeing the write buffers in the error path of the
    ubifs_remount_rw() is wrong. It leads later to a kernel oops like this:

    [10016.431274] UBIFS (ubi0:0): start fixing up free space
    [10090.810042] UBIFS (ubi0:0): free space fixup complete
    [10090.814623] UBIFS error (ubi0:0 pid 512): ubifs_remount_fs: cannot
    spawn "ubifs_bgt0_0", error -4
    [10101.915108] UBIFS (ubi0:0): background thread "ubifs_bgt0_0" started,
    PID 517
    [10105.275498] Unable to handle kernel NULL pointer dereference at
    virtual address 0000000000000030
```

```
[10105.284352] Mem abort info:
[10105.287160]   ESR = 0x96000006
[10105.290252]   EC = 0x25: DABT (current EL), IL = 32 bits
[10105.295592]   SET = 0, FnV = 0
[10105.298652]   EA = 0, S1PTW = 0
[10105.301848] Data abort info:
[10105.304723]   ISV = 0, ISS = 0x00000006
[10105.308573]   CM = 0, WnR = 0
[10105.311564] user pgtable: 4k pages, 48-bit VAs, pgdp=00000000f03d1000
[10105.318034] [0000000000000030] pgd=00000000f6cee003,
pud=00000000f4884003, pmd=0000000000000000
[10105.326783] Internal error: Oops: 96000006 [#1] PREEMPT SMP
[10105.332355] Modules linked in: ath10k_pci ath10k_core ath mac80211
libarc4 cfg80211 nvme nvme_core cryptodev(O)
[10105.342468] CPU: 3 PID: 518 Comm: touch Tainted: G           O
5.4.3 #1
[10105.349517] Hardware name: HYPEX CPU (DT)
[10105.353525] pstate: 40000005 (nZcv daif -PAN -UAO)
[10105.358324] pc : atomic64_try_cmpxchg_acquire.constprop.22+0x8/0x34
[10105.364596] lr : mutex_lock+0x1c/0x34
[10105.368253] sp : ffff000075633aa0
[10105.371563] x29: ffff000075633aa0 x28: 0000000000000001
[10105.376874] x27: ffff000076fa80c8 x26: 0000000000000004
[10105.382185] x25: 0000000000000030 x24: 0000000000000000
[10105.387495] x23: 0000000000000000 x22: 0000000000000038
[10105.392807] x21: 000000000000000c x20: ffff000076fa80c8
[10105.398119] x19: ffff000076fa8000 x18: 0000000000000000
[10105.403429] x17: 0000000000000000 x16: 0000000000000000
[10105.408741] x15: 0000000000000000 x14: fefefefefefefeff
[10105.414052] x13: 0000000000000000 x12: 0000000000000fe0
[10105.419364] x11: 0000000000000fe0 x10: ffff000076709020
[10105.424675] x9 : 0000000000000000 x8 : 00000000000000a0
[10105.429986] x7 : ffff000076fa80f4 x6 : 0000000000000030
[10105.435297] x5 : 0000000000000000 x4 : 0000000000000000
[10105.440609] x3 : 0000000000000000 x2 : ffff00006f276040
[10105.445920] x1 : ffff000075633ab8 x0 : 0000000000000030
[10105.451232] Call trace:
[10105.453676]  atomic64_try_cmpxchg_acquire.constprop.22+0x8/0x34
[10105.459600]  ubifs_garbage_collect+0xb4/0x334
[10105.463956]  ubifs_budget_space+0x398/0x458
[10105.468139]  ubifs_create+0x50/0x180
[10105.471712]  path_openat+0x6a0/0x9b0
[10105.475284]  do_filp_open+0x34/0x7c
[10105.478771]  do_sys_open+0x78/0xe4
[10105.482170]  __arm64_sys_openat+0x1c/0x24
[10105.486180]  el0_svc_handler+0x84/0xc8
[10105.489928]  el0_svc+0x8/0xc
[10105.492808] Code: 52800013 17fffffb d2800003 f9800011 (c85ffc05)
[10105.498903] ---[ end trace 46b721d93267a586 ]---

To reproduce the problem:

1. Filesystem initially mounted read-only, free space fixup flag set.

2. mount -o remount,rw <mountpoint>

3. it takes some time (free space fixup running)
    ... try to terminate running mount by CTRL-C
    ... does not respond, only after free space fixup is complete
    ... then "ubifs_remount_fs: cannot spawn "ubifs_bgt0_0", error -4"

4. mount -o remount,rw <mountpoint>
    ... now finished instantly (fixup already done).

5. Create file or just unmount the filesystem and we get the oops.
```

```
    Cc: <stable@vger.kernel.org>
    Fixes: b50b9f408502 ("UBIFS: do not free write-buffers when in R/O mode")
    Signed-off-by: Petr Cvachoucek <cvachoucek@gmail.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ecf0d470d1abdb90448026fbee6b63069ec8ce1b
Author: Meng Li <Meng.Li@windriver.com>
Date:   Mon Nov 1 11:13:53 2021 +0800

    crypto: caam - replace this_cpu_ptr with raw_cpu_ptr

    commit efd21e10fc3bf4c6da122470a5ae89ec4ed8d180 upstream.

    When enable the kernel debug config, there is below calltrace detected:
    BUG: using smp_processor_id() in preemptible [00000000] code: cryptomgr_test/339
    caller is debug_smp_processor_id+0x20/0x30
    CPU: 9 PID: 339 Comm: cryptomgr_test Not tainted 5.10.63-yocto-standard #1
    Hardware name: NXP Layerscape LX2160ARDB (DT)
    Call trace:
     dump_backtrace+0x0/0x1a0
     show_stack+0x24/0x30
     dump_stack+0xf0/0x13c
     check_preemption_disabled+0x100/0x110
     debug_smp_processor_id+0x20/0x30
     dpaa2_caam_enqueue+0x10c/0x25c
     ......
     cryptomgr_test+0x38/0x60
     kthread+0x158/0x164
     ret_from_fork+0x10/0x38
    According to the comment in commit ac5d15b4519f("crypto: caam/qi2
     - use affine DPIOs "), because preemption is no longer disabled
    while trying to enqueue an FQID, it might be possible to run the
    enqueue on a different CPU(due to migration, when in process context),
    however this wouldn't be a functionality issue. But there will be
    above calltrace when enable kernel debug config. So, replace this_cpu_ptr
    with raw_cpu_ptr to avoid above call trace.

    Fixes: ac5d15b4519f ("crypto: caam/qi2 - use affine DPIOs")
    Cc: stable@vger.kernel.org
    Signed-off-by: Meng Li <Meng.Li@windriver.com>
    Reviewed-by: Horia Geantă <horia.geanta@nxp.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3536c2e8958b4b759663a0f103534287b15b57dc
Author: Marek Vasut <marex@denx.de>
Date:   Mon Dec 20 20:50:22 2021 +0100

    crypto: stm32/crc32 - Fix kernel BUG triggered in probe()

    commit 29009604ad4e3ef784fd9b9fef6f23610ddf633d upstream.

    The include/linux/crypto.h struct crypto_alg field cra_driver_name description
    states "Unique name of the transformation provider. " ... " this contains the
    name of the chip or provider and the name of the transformation algorithm."

    In case of the stm32-crc driver, field cra_driver_name is identical for all
    registered transformation providers and set to the name of the driver itself,
    which is incorrect. This patch fixes it by assigning a unique cra_driver_name
    to each registered transformation provider.

    The kernel crash is triggered when the driver calls crypto_register_shashes()
    which calls crypto_register_shash(), which calls crypto_register_alg(), which
```

```
    calls __crypto_register_alg(), which returns -EEXIST, which is propagated
    back through this call chain. Upon -EEXIST from crypto_register_shash(), the
    crypto_register_shashes() starts unregistering the providers back, and calls
    crypto_unregister_shash(), which calls crypto_unregister_alg(), and this is
    where the BUG() triggers due to incorrect cra_refcnt.

    Fixes: b51dbe90912a ("crypto: stm32 - Support for STM32 CRC32 crypto module")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Cc: <stable@vger.kernel.org> # 4.12+
    Cc: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Cc: Fabien Dessenne <fabien.dessenne@st.com>
    Cc: Herbert Xu <herbert@gondor.apana.org.au>
    Cc: Lionel Debieve <lionel.debieve@st.com>
    Cc: Nicolas Toromanoff <nicolas.toromanoff@st.com>
    Cc: linux-arm-kernel@lists.infradead.org
    Cc: linux-stm32@st-md-mailman.stormreply.com
    To: linux-crypto@vger.kernel.org
    Acked-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 967d28f8128f45a1f4cdc229fec17f02bd15167f
Author: Heiner Kallweit <hkallweit1@gmail.com>
Date:   Fri Dec 17 10:03:30 2021 +0100

    crypto: omap-aes - Fix broken pm_runtime_and_get() usage

    commit c2aec59be093bd44627bc4f6bc67e4614a93a7b6 upstream.

    This fix is basically the same as 3d6b661330a7 ("crypto: stm32 -
    Revert broken pm_runtime_resume_and_get changes"), just for the omap
    driver. If the return value isn't used, then pm_runtime_get_sync()
    has to be used for ensuring that the usage count is balanced.

    Fixes: 1f34cc4a8da3 ("crypto: omap-aes - Fix PM reference leak on omap-aes.c")
    Cc: stable@vger.kernel.org
    Signed-off-by: Heiner Kallweit <hkallweit1@gmail.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e09b9217cbecf318b4580cecdd6ab531d5266b8d
Author: Zhu Lingshan <lingshan.zhu@intel.com>
Date:   Wed Dec 1 16:12:55 2021 +0800

    ifcvf/vDPA: fix misuse virtio-net device config size for blk dev

    commit 0f420c383a2bb414ebccedf9289b5b815f1295fe upstream.

    This commit fixes a misuse of virtio-net device config size issue
    for virtio-block devices.

    A new member config_size in struct ifcvf_hw is introduced and would
    be initialized through vdpa_dev_add() to record correct device
    config size.

    To be more generic, rename ifcvf_hw.net_config to ifcvf_hw.dev_config,
    the helpers ifcvf_read/write_net_config() to ifcvf_read/write_dev_config()

    Signed-off-by: Zhu Lingshan <lingshan.zhu@intel.com>
    Reported-and-suggested-by: Stefano Garzarella <sgarzare@redhat.com>
    Reviewed-by: Stefano Garzarella <sgarzare@redhat.com>
    Fixes: 6ad31d162a4e ("vDPA/ifcvf: enable Intel C5000X-PL virtio-block for vDPA")
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20211201081255.60187-1-lingshan.zhu@intel.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
```

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

commit 75bb9241e68d38061c7d7c9a0091dd410f51dd88
Author: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
Date:   Mon Dec 6 20:07:58 2021 +0100

    rpmsg: core: Clean up resources on announce_create failure.

    commit 8066c615cb69b7da8a94f59379847b037b3a5e46 upstream.

    During the rpmsg_dev_probe, if rpdev->ops->announce_create returns an
    error, the rpmsg device and default endpoint should be freed before
    exiting the function.

    Fixes: 5e619b48677c ("rpmsg: Split rpmsg core and virtio backend")
    Suggested-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20211206190758.10004-1-arnaud.pouliquen@foss.st.com
    Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0c22cec4550fdff80766ebd817fab4227de670c0
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Jan 14 14:08:30 2022 -0800

    hugetlbfs: fix off-by-one error in hugetlb_vmdelete_list()

    [ Upstream commit d6aba4c8e20d4d2bf65d589953f6d891c178f3a3 ]

    Pass "end - 1" instead of "end" when walking the interval tree in
    hugetlb_vmdelete_list() to fix an inclusive vs.  exclusive bug.  The two
    callers that pass a non-zero "end" treat it as exclusive, whereas the
    interval tree iterator expects an inclusive "last".  E.g.  punching a
    hole in a file that precisely matches the size of a single hugepage,
    with a vma starting right on the boundary, will result in
    unmap_hugepage_range() being called twice, with the second call having
    start==end.

    The off-by-one error doesn't cause functional problems as
    __unmap_hugepage_range() turns into a massive nop due to
    short-circuiting its for-loop on "address < end".  But, the mmu_notifier
    invocations to invalid_range_{start,end}() are passed a bogus zero-sized
    range, which may be unexpected behavior for secondary MMUs.

    The bug was exposed by commit ed922739c919 ("KVM: Use interval tree to
    do fast hva lookup in memslots"), currently queued in the KVM tree for
    5.17, which added a WARN to detect ranges with start==end.

    Link: https://lkml.kernel.org/r/20211228234257.1926057-1-seanjc@google.com
    Fixes: 1bfad99ab425 ("hugetlbfs: hugetlb_vmtruncate_list() needs to take a range to
delete")
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Reported-by: syzbot+4e697fe80a31aa7efe21@syzkaller.appspotmail.com
    Reviewed-by: Mike Kravetz <mike.kravetz@oracle.com>
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a139bc4cf9ffe905913deee3f88e16876828155
Author: Waiman Long <longman@redhat.com>
Date:   Fri Jan 14 14:07:58 2022 -0800
```

```
    selftests/vm: make charge_reserved_hugetlb.sh work with existing cgroup setting

    [ Upstream commit 209376ed2a8431ccb4c40fdcef11194fc1e749b0 ]

    The hugetlb cgroup reservation test charge_reserved_hugetlb.sh assume
    that no cgroup filesystems are mounted before running the test.  That is
    not true in many cases.  As a result, the test fails to run.  Fix that
    by querying the current cgroup mount setting and using the existing
    cgroup setup instead before attempting to freshly mount a cgroup
    filesystem.

    Similar change is also made for hugetlb_reparenting_test.sh as well,
    though it still has problem if cgroup v2 isn't used.

    The patched test scripts were run on a centos 8 based system to verify
    that they ran properly.

    Link: https://lkml.kernel.org/r/20220106201359.1646575-1-longman@redhat.com
    Fixes: 29750f71a9b4 ("hugetlb_cgroup: add hugetlb_cgroup reservation tests")
    Signed-off-by: Waiman Long <longman@redhat.com>
    Acked-by: Mina Almasry <almasrymina@google.com>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: Mike Kravetz <mike.kravetz@oracle.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fab2f1c3141bc5550b5ec376595ae9a091cdd657
Author: Andrey Konovalov <andreyknvl@gmail.com>
Date:   Fri Jan 14 14:05:01 2022 -0800

    kasan: fix quarantine conflicting with init_on_free

    [ Upstream commit 26dca996ea7b1ac7008b6b6063fc88b849e3ac3e ]

    KASAN's quarantine might save its metadata inside freed objects.  As
    this happens after the memory is zeroed by the slab allocator when
    init_on_free is enabled, the memory coming out of quarantine is not
    properly zeroed.

    This causes lib/test_meminit.c tests to fail with Generic KASAN.

    Zero the metadata when the object is removed from quarantine.

    Link:
https://lkml.kernel.org/r/2805da5df4b57138fdacd671f5d227d58950ba54.1640037083.git.andreyknvl
@google.com
    Fixes: 6471384af2a6 ("mm: security: introduce init_on_alloc=1 and init_on_free=1 boot
options")
    Signed-off-by: Andrey Konovalov <andreyknvl@google.com>
    Reviewed-by: Marco Elver <elver@google.com>
    Cc: Alexander Potapenko <glider@google.com>
    Cc: Andrey Konovalov <andreyknvl@gmail.com>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: Andrey Ryabinin <ryabinin.a.a@gmail.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 218ceb8b5ca8e7cece8c7872161649870e0ae5da
Author: Kefeng Wang <wangkefeng.wang@huawei.com>
Date:   Fri Jan 14 14:04:11 2022 -0800

    mm: defer kmemleak object creation of module_alloc()
```

```
    [ Upstream commit 60115fa54ad7b913b7cb5844e6b7ffeb842d55f2 ]

    Yongqiang reports a kmemleak panic when module insmod/rmmod with KASAN
    enabled(without KASAN_VMALLOC) on x86[1].

    When the module area allocates memory, it's kmemleak_object is created
    successfully, but the KASAN shadow memory of module allocation is not
    ready, so when kmemleak scan the module's pointer, it will panic due to
    no shadow memory with KASAN check.

      module_alloc
        __vmalloc_node_range
          kmemleak_vmalloc
                                  kmemleak_scan
                                    update_checksum
        kasan_module_alloc
          kmemleak_ignore

    Note, there is no problem if KASAN_VMALLOC enabled, the modules area
    entire shadow memory is preallocated.  Thus, the bug only exits on ARCH
    which supports dynamic allocation of module area per module load, for
    now, only x86/arm64/s390 are involved.

    Add a VM_DEFER_KMEMLEAK flags, defer vmalloc'ed object register of
    kmemleak in module_alloc() to fix this issue.

    [1] https://lore.kernel.org/all/6d41e2b9-4692-5ec4-b1cd-cbe29ae89739@huawei.com/

    [wangkefeng.wang@huawei.com: fix build]
      Link: https://lkml.kernel.org/r/20211125080307.27225-1-wangkefeng.wang@huawei.com
    [akpm@linux-foundation.org: simplify ifdefs, per Andrey]
      Link: https://lkml.kernel.org/r/CA+fCnZcnwJHUQq34VuRxpdoY6_XbJCDJ-
jopksS5Eia4PijPzw@mail.gmail.com

    Link: https://lkml.kernel.org/r/20211124142034.192078-1-wangkefeng.wang@huawei.com
    Fixes: 793213a82de4 ("s390/kasan: dynamic shadow mem allocation for modules")
    Fixes: 39d114ddc682 ("arm64: add KASAN support")
    Fixes: bebf56a1b176 ("kasan: enable instrumentation of global variables")
    Signed-off-by: Kefeng Wang <wangkefeng.wang@huawei.com>
    Reported-by: Yongqiang Liu <liuyongqiang13@huawei.com>
    Cc: Andrey Konovalov <andreyknvl@gmail.com>
    Cc: Andrey Ryabinin <ryabinin.a.a@gmail.com>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Will Deacon <will@kernel.org>
    Cc: Heiko Carstens <hca@linux.ibm.com>
    Cc: Vasily Gorbik <gor@linux.ibm.com>
    Cc: Christian Borntraeger <borntraeger@linux.ibm.com>
    Cc: Alexander Gordeev <agordeev@linux.ibm.com>
    Cc: Thomas Gleixner <tglx@linutronix.de>
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Borislav Petkov <bp@alien8.de>
    Cc: Dave Hansen <dave.hansen@linux.intel.com>
    Cc: Alexander Potapenko <glider@google.com>
    Cc: Kefeng Wang <wangkefeng.wang@huawei.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b1e0048d728a36d42a042603a8d39d7b61c2e386
  Author: Xiaoke Wang <xkernel.wang@foxmail.com>
  Date:   Tue Dec 14 10:26:46 2021 +0800

    tracing/probes: check the return value of kstrndup() for pbuf
```

```
    [ Upstream commit 1c1857d400355e96f0fe8b32adc6fa7594d03b52 ]

    kstrndup() is a memory allocation-related function, it returns NULL when
    some internal memory errors happen. It is better to check the return
    value of it so to catch the memory error in time.

    Link: https://lkml.kernel.org/r/tencent_4D6E270731456EB88712ED7F13883C334906@qq.com

    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Fixes: a42e3c4de964 ("tracing/probe: Add immediate string parameter support")
    Signed-off-by: Xiaoke Wang <xkernel.wang@foxmail.com>
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06e2dedfddcc25b6db3d00db064e910aef2154e2
Author: Xiaoke Wang <xkernel.wang@foxmail.com>
Date:   Tue Dec 14 09:28:02 2021 +0800

    tracing/uprobes: Check the return value of kstrdup() for tu->filename

    [ Upstream commit 8c7224245557707c613f130431cafbaaa4889615 ]

    kstrdup() returns NULL when some internal memory errors happen, it is
    better to check the return value of it so to catch the memory error in
    time.

    Link: https://lkml.kernel.org/r/tencent_3C2E330722056D7891D2C83F29C802734B06@qq.com

    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Fixes: 33ea4b24277b ("perf/core: Implement the 'perf_uprobe' PMU")
    Signed-off-by: Xiaoke Wang <xkernel.wang@foxmail.com>
    Signed-off-by: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aadeb5dd1a22125495eef4c1d7bdb84bb8fbbb78
Author: Weizhao Ouyang <o451686892@gmail.com>
Date:   Tue Jan 4 15:35:45 2022 +0800

    dma-buf: cma_heap: Fix mutex locking section

    [ Upstream commit 54329e6f7beea6af56c1230da293acc97d6a6ee7 ]

    Fix cma_heap_buffer mutex locking critical section to protect vmap_cnt
    and vaddr.

    Fixes: a5d2d29e24be ("dma-buf: heaps: Move heap-helper logic into the cma_heap
implementation")
    Signed-off-by: Weizhao Ouyang <o451686892@gmail.com>
    Acked-by: John Stultz <john.stultz@linaro.org>
    Signed-off-by: Sumit Semwal <sumit.semwal@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220104073545.124244-1-
o451686892@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d477d83659960f7b59a824b3108ef450b92844b5
Author: Tom Rix <trix@redhat.com>
Date:   Sat Jan 8 07:09:48 2022 -0800

    i3c: master: dw: check return of dw_i3c_master_get_free_pos()

    [ Upstream commit 13462ba1815db5a96891293a9cfaa2451f7bd623 ]

    Clang static analysis reports this problem
    dw-i3c-master.c:799:9: warning: The result of the left shift is
      undefined because the left operand is negative
```

```
                            COMMAND_PORT_DEV_INDEX(pos) |
                            ^~~~~~~~~~~~~~~~~~~~~~~~~~~
```

    pos can be negative because dw_i3c_master_get_free_pos() can return an
    error. So check for an error.

    Fixes: 1dd728f5d4d4 ("i3c: master: Add driver for Synopsys DesignWare IP")
    Signed-off-by: Tom Rix <trix@redhat.com>
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/20220108150948.3988790-1-trix@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0181c66114734cb0c0d858e1ebdd46dac2199f2a
Author: Sergio Paracuellos <sergio.paracuellos@gmail.com>
Date: Tue Dec 7 11:49:23 2021 +0100

    PCI: mt7621: Add missing MODULE_LICENSE()

    [Upstream commit e4b1cd02dc8d7967a79edccd510724831e5cdee8]

    The MT7621 PCIe host controller driver can be built as a module, but it
    lacks a MODULE_LICENSE(), which causes a build error:

      ERROR: modpost: missing MODULE_LICENSE() in drivers/pci/controller/pcie-mt7621.o

    Add MODULE_LICENSE() to the driver.

    Fixes: 2bdd5238e756 ("PCI: mt7621: Add MediaTek MT7621 PCIe host controller driver")
    Link: https://lore.kernel.org/r/20211207104924.21327-5-sergio.paracuellos@gmail.com
    Signed-off-by: Yanteng Si <siyanteng@loongson.cn>
    Signed-off-by: Sergio Paracuellos <sergio.paracuellos@gmail.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Krzysztof Wilczyński <kw@linux.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e684eb8e36a6cc829c53868c269ef06abe534f3
Author: Guchun Chen <guchun.chen@amd.com>
Date: Fri Jan 7 16:31:20 2022 +0800

    drm/amdgpu: use spin_lock_irqsave to avoid deadlock by local interrupt

    [Upstream commit 2096b74b1da5ca418827b54ac4904493bd9de89c]

    This is observed in SRIOV case with virtual KMS as display.

    _raw_spin_lock_irqsave+0x37/0x40
    drm_handle_vblank+0x69/0x350 [drm]
    ? try_to_wake_up+0x432/0x5c0
    ? amdgpu_vkms_prepare_fb+0x1c0/0x1c0 [amdgpu]
    drm_crtc_handle_vblank+0x17/0x20 [drm]
    amdgpu_vkms_vblank_simulate+0x4d/0x80 [amdgpu]
    __hrtimer_run_queues+0xfb/0x230
    hrtimer_interrupt+0x109/0x220
    __sysvec_apic_timer_interrupt+0x64/0xe0
    asm_call_irq_on_stack+0x12/0x20

    Fixes: 84ec374bd580 ("drm/amdgpu: create amdgpu_vkms (v4)")
    Signed-off-by: Guchun Chen <guchun.chen@amd.com>
    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Tested-by: Kelly Zytaruk <kelly.zytaruk@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 506c9632d77c0ae755fb66f5a0b8578c0b65a84b
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>

```
Date:   Wed Jan 5 17:09:43 2022 +0800

    drm/amdkfd: Check for null pointer after calling kmemdup

    [ Upstream commit abfaf0eee97925905e742aa3b0b72e04a918fa9e ]

    As the possible failure of the allocation, kmemdup() may return NULL
    pointer.
    Therefore, it should be better to check the 'props2' in order to prevent
    the dereference of NULL pointer.

    Fixes: 3a87177eb141 ("drm/amdkfd: Add topology support for dGPUs")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Felix Kuehling <Felix.Kuehling@amd.com>
    Signed-off-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 d42642b662e740b772cd1a5ef011234cc6172f97
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Sun Jan 9 19:42:45 2022 +0100

    drm/amd/display: invalid parameter check in dmub_hpd_callback

    [ Upstream commit 978ffac878fd64039f95798b15b430032d2d89d5 ]

    The function performs a check on the "adev" input parameter, however, it
    is used before the check.

    Initialize the "dev" variable after the sanity check to avoid a possible
    NULL pointer dereference.

    Fixes: e27c41d5b0681 ("drm/amd/display: Support for DMUB HPD interrupt handling")
    Addresses-Coverity-ID: 1493909 ("Null pointer dereference")
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2c6766ab56cb41c27ca60101dfc857496d65ceee
Author: Wesley Sheng <wesley.sheng@microchip.com>
Date:   Thu Dec 23 17:23:30 2021 -0800

    ntb_hw_switchtec: Fix bug with more than 32 partitions

    [ Upstream commit 7ff351c86b6b258f387502ab2c9b9d04f82c1c3d ]

    Switchtec could support as mush as 48 partitions, but ffs & fls are
    for 32 bit argument, in case of partition index larger than 31, the
    current code could not parse the peer partition index correctly.
    Change to the 64 bit version __ffs64 & fls64 accordingly to fix this
    bug.

    Fixes: 3df54c870f52 ("ntb_hw_switchtec: Allow using Switchtec NTB in multi-partition
setups")
    Signed-off-by: Wesley Sheng <wesley.sheng@microchip.com>
    Signed-off-by: Kelvin Cao <kelvin.cao@microchip.com>
    Signed-off-by: Jon Mason <jdmason@kudzu.us>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bb9709986133465c3cfb71f348250982ebc2ddc0
Author: Jeremy Pallotta <jmpallotta@gmail.com>
Date:   Thu Dec 23 17:23:29 2021 -0800

    ntb_hw_switchtec: Fix pff ioread to read into mmio_part_cfg_all
```

```
    [ Upstream commit 32c3d375b0ed84b6acb51ae5ebef35ff0d649d85 ]

    Array mmio_part_cfg_all holds the partition configuration of all
    partitions, with partition number as index. Fix this by reading into
    mmio_part_cfg_all for pff.

    Fixes: 0ee28f26f378 ("NTB: switchtec_ntb: Add link management")
    Signed-off-by: Jeremy Pallotta <jmpallotta@gmail.com>
    Signed-off-by: Kelvin Cao <kelvin.cao@microchip.com>
    Signed-off-by: Jon Mason <jdmason@kudzu.us>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b929afe3a8ebc86dcbe7fbd2bc1d97d23e81731e
Author: Liu Ying <victor.liu@nxp.com>
Date:   Thu Dec 30 12:06:26 2021 +0800

    drm/atomic: Check new_crtc_state->active to determine if CRTC needs disable in self
refresh mode

    [ Upstream commit 69e630016ef4e4a1745310c446f204dc6243e907 ]

    Actual hardware state of CRTC is controlled by the member 'active' in
    struct drm_crtc_state instead of the member 'enable', according to the
    kernel doc of the member 'enable'.  In fact, the drm client modeset
    and atomic helpers are using the member 'active' to do the control.

    Referencing the member 'enable' of new_crtc_state, the function
    crtc_needs_disable() may fail to reflect if CRTC needs disable in
    self refresh mode, e.g., when the framebuffer emulation will be blanked
    through the client modeset helper with the next commit, the member
    'enable' of new_crtc_state is still true while the member 'active' is
    false, hence the relevant potential encoder and bridges won't be disabled.

    So, let's check new_crtc_state->active to determine if CRTC needs disable
    in self refresh mode instead of new_crtc_state->enable.

    Fixes: 1452c25b0e60 ("drm: Add helpers to kick off self refresh mode in drivers")
    Cc: Sean Paul <seanpaul@chromium.org>
    Cc: Rob Clark <robdclark@chromium.org>
    Cc: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    Cc: Maxime Ripard <mripard@kernel.org>
    Cc: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: David Airlie <airlied@linux.ie>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Liu Ying <victor.liu@nxp.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211230040626.646807-1-
victor.liu@nxp.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a570c79284bbabcf0658cbef8139ca408e279b4d
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Jan 7 08:36:32 2022 +0000

    drm/sun4i: dw-hdmi: Fix missing put_device() call in sun8i_hdmi_phy_get

    [ Upstream commit c71af3dae3e34d2fde0c19623cf7f8483321f0e3 ]

    The reference taken by 'of_find_device_by_node()' must be released when
    not needed anymore.
    Add the corresponding 'put_device()' in the error handling path.

    Fixes: 9bf3797796f5 ("drm/sun4i: dw-hdmi: Make HDMI PHY into a platform device")
```

```
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220107083633.20843-1-
linmq006@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3410a321d10e589536c7fd71d9bb354830691405
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Sat Jan 8 16:59:54 2022 -0500

    SUNRPC: Fix sockaddr handling in svcsock_accept_class trace points

    [ Upstream commit 16720861675393a35974532b3c837d9fd7bfe08c ]

    Avoid potentially hazardous memory copying and the needless use of
    "%pIS" -- in the kernel, an RPC service listener is always bound to
    ANYADDR. Having the network namespace is helpful when recording
    errors, though.

    Fixes: a0469f46faab ("SUNRPC: Replace dprintk call sites in TCP state change callouts")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0cfc5a11bbb360bfdb58d983d587ae818a7e70a5
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Sun Jan 9 13:26:51 2022 -0500

    SUNRPC: Fix sockaddr handling in the svc_xprt_create_error trace point

    [ Upstream commit dc6c6fb3d639756a532bcc47d4a9bf9f3965881b ]

    While testing, I got an unexpected KASAN splat:

    Jan 08 13:50:27 oracle-102.nfsv4.dev kernel: BUG: KASAN: stack-out-of-bounds in
trace_event_raw_event_svc_xprt_create_err+0x190/0x210 [sunrpc]
    Jan 08 13:50:27 oracle-102.nfsv4.dev kernel: Read of size 28 at addr ffffc9000008f728 by
task mount.nfs/4628

    The memcpy() in the TP_fast_assign section of this trace point
    copies the size of the destination buffer in order that the buffer
    won't be overrun.

    In other similar trace points, the source buffer for this memcpy is
    a "struct sockaddr_storage" so the actual length of the source
    buffer is always long enough to prevent the memcpy from reading
    uninitialized or unallocated memory.

    However, for this trace point, the source buffer can be as small as
    a "struct sockaddr_in". For AF_INET sockaddrs, the memcpy() reads
    memory that follows the source buffer, which is not always valid
    memory.

    To avoid copying past the end of the passed-in sockaddr, make the
    source address's length available to the memcpy(). It would be a
    little nicer if the tracing infrastructure was more friendly about
    storing socket addresses that are not AF_INET, but I could not find
    a way to make printk("%pIS") work with a dynamic array.

    Reported-by: KASAN
    Fixes: 4b8f380e46e4 ("SUNRPC: Tracepoint to record errors in svc_xpo_create()")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba806da3b75d7361381bcfa87dd27593baf97d07
Author: Matthew Auld <matthew.auld@intel.com>
```

```
Date:   Thu Jan 6 17:49:07 2022 +0000

    drm/i915: don't call free_mmap_offset when purging

    [ Upstream commit 4c2602ba8d74c35d550ed3d518809c697de08d88 ]

    The TTM backend is in theory the only user here(also purge should only
    be called once we have dropped the pages), where it is setup at object
    creation and is only removed once the object is destroyed. Also
    resetting the node here might be iffy since the ttm fault handler
    uses the stored fake offset to determine the page offset within the pages
    array.

    This also blows up in the dontneed-before-mmap test, since the
    expectation is that the vma_node will live on, until the object is
    destroyed:

    <2> [749.062902] kernel BUG at drivers/gpu/drm/i915/gem/i915_gem_ttm.c:943!
    <4> [749.062923] invalid opcode: 0000 [#1] PREEMPT SMP NOPTI
    <4> [749.062928] CPU: 0 PID: 1643 Comm: gem_madvise Tainted: G     U  W         5.16.0-
rc8-CI-CI_DRM_11046+ #1
    <4> [749.062933] Hardware name: Gigabyte Technology Co., Ltd. GB-Z390 Garuda/GB-Z390
Garuda-CF, BIOS IG1c 11/19/2019
    <4> [749.062937] RIP: 0010:i915_ttm_mmap_offset.cold.35+0x5b/0x5d [i915]
    <4> [749.063044] Code: 00 48 c7 c2 a0 23 4e a0 48 c7 c7 26 df 4a a0 e8 95 1d d0 e0 bf 01
00 00 00 e8 8b ec cf e0 31 f6 bf 09 00 00 00 e8 5f 30 c0 e0 <0f> 0b 48 c7 c1 24 4b 56 a0 ba
5b 03 00 00 48 c7 c6 c0 23 4e a0 48
    <4> [749.063052] RSP: 0018:ffffc90002ab7d38 EFLAGS: 00010246
    <4> [749.063056] RAX: 0000000000000240 RBX: ffff88811f2e61c0 RCX: 0000000000000006
    <4> [749.063060] RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000009
    <4> [749.063063] RBP: ffffc90002ab7e58 R08: 0000000000000001 R09: 0000000000000001
    <4> [749.063067] R10: 000000000123d0f8 R11: ffffc90002ab7b20 R12: ffff888112a1a000
    <4> [749.063071] R13: 0000000000000004 R14: ffff88811f2e61c0 R15: ffff888112a1a000
    <4> [749.063074] FS:  00007f6e5fcad500(0000) GS:ffff8884ad600000(0000)
knlGS:0000000000000000
    <4> [749.063078] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    <4> [749.063081] CR2: 00007efd264e39f0 CR3: 0000000115fd6005 CR4: 00000000003706f0
    <4> [749.063085] Call Trace:
    <4> [749.063087]  <TASK>
    <4> [749.063089]  __assign_mmap_offset+0x41/0x300 [i915]
    <4> [749.063171]  __assign_mmap_offset_handle+0x159/0x270 [i915]
    <4> [749.063248]  ? i915_gem_dumb_mmap_offset+0x70/0x70 [i915]
    <4> [749.063325]  drm_ioctl_kernel+0xae/0x140
    <4> [749.063330]  drm_ioctl+0x201/0x3d0
    <4> [749.063333]  ? i915_gem_dumb_mmap_offset+0x70/0x70 [i915]
    <4> [749.063409]  ? do_user_addr_fault+0x200/0x670
    <4> [749.063415]  __x64_sys_ioctl+0x6d/0xa0
    <4> [749.063419]  do_syscall_64+0x3a/0xb0
    <4> [749.063423]  entry_SYSCALL_64_after_hwframe+0x44/0xae
    <4> [749.063428] RIP: 0033:0x7f6e5f100317

    Testcase: igt/gem_madvise/dontneed-before-mmap
    Fixes: cf3e3e86d779 ("drm/i915: Use ttm mmap handling for ttm bo's.")
    Signed-off-by: Matthew Auld <matthew.auld@intel.com>
    Cc: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Reviewed-by: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220106174910.280616-1-
matthew.auld@intel.com
    (cherry picked from commit 658a0c632625e1db51837ff754fe18a6a7f2ccf8)
    Signed-off-by: Tvrtko Ursulin <tvrtko.ursulin@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 089042e7d524f4d3fabb030ab44f1dc70b5e762e
Author: Juston Li <juston.li@intel.com>
Date:   Thu Jan 6 12:02:36 2022 -0800
```

```
drm/i915/pxp: Hold RPM wakelock during PXP unbind

[ Upstream commit f9535d28ac93c3cc326f7215fccd0abe1d3a6083 ]

Similar to commit b8d8436840ca ("drm/i915/gt: Hold RPM wakelock during
PXP suspend") but to fix the same warning for unbind during shutdown:

------------[ cut here ]------------
RPM wakelock ref not held during HW access
WARNING: CPU: 0 PID: 4139 at drivers/gpu/drm/i915/intel_runtime_pm.h:115
gen12_fwtable_write32+0x1b7/0
Modules linked in: 8021q ccm rfcomm cmac algif_hash algif_skcipher
af_alg uinput snd_hda_codec_hdmi vf industrialio iwl7000_mac80211
cros_ec_sensorhub lzo_rle lzo_compress zram iwlwifi cfg80211 joydev
CPU: 0 PID: 4139 Comm: halt Tainted: G     U  W
5.10.84 #13 344e11e079c4a03940d949e537eab645f6
RIP: 0010:gen12_fwtable_write32+0x1b7/0x200
Code: 48 c7 c7 fc b3 b5 89 31 c0 e8 2c f3 ad ff 0f 0b e9 04 ff ff ff c6
05 71 e9 1d 01 01 48 c7 c7 d67
RSP: 0018:ffffa09ec0bb3bb0 EFLAGS: 00010246
RAX: 12dde97bbd260300 RBX: 00000000000320f0 RCX: ffffffff89e60ea0
RDX: 0000000000000000 RSI: 00000000ffffdfff RDI: ffffffff89e60e70
RBP: ffffa09ec0bb3bd8 R08: 0000000000000000 R09: ffffa09ec0bb3950
R10: 00000000ffffdfff R11: ffffffff89e91160 R12: 0000000000000000
R13: 0000000028121969 R14: ffff9515c32f0990 R15: 0000000040000000
FS:  0000790dcf225740(0000) GS:ffff951737800000(0000) knlGS:0000000000000000
CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
CR2: 000058b25efae147 CR3: 0000000133ea6001 CR4: 0000000000770ef0
DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
DR3: 0000000000000000 DR6: 00000000ffff07f0 DR7: 0000000000000400
PKRU: 55555554
Call Trace:
 intel_pxp_fini_hw+0x2f/0x39
 i915_pxp_tee_component_unbind+0x1c/0x42
 component_unbind+0x32/0x48
 component_unbind_all+0x80/0x9d
 take_down_master+0x24/0x36
 component_master_del+0x56/0x70
 mei_pxp_remove+0x2c/0x68
 mei_cl_device_remove+0x35/0x68
 device_release_driver_internal+0x100/0x1a1
 mei_cl_bus_remove_device+0x21/0x79
 mei_cl_bus_remove_devices+0x3b/0x51
 mei_stop+0x3b/0xae
 mei_me_shutdown+0x23/0x58
 device_shutdown+0x144/0x1d3
 kernel_power_off+0x13/0x4c
 __se_sys_reboot+0x1d4/0x1e9
 do_syscall_64+0x43/0x55
 entry_SYSCALL_64_after_hwframe+0x44/0xa9
RIP: 0033:0x790dcf316273
Code: 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 44 00
00 89 fa be 69 19 12 28 bf ad8
RSP: 002b:00007ffca0df9198 EFLAGS: 00000202 ORIG_RAX: 00000000000000a9
RAX: ffffffffffffffda RBX: 000000004321fedc RCX: 0000790dcf316273
RDX: 000000004321fedc RSI: 0000000028121969 RDI: 00000000fee1dead
RBP: 00007ffca0df9200 R08: 0000000000000007 R09: 0000563ce8cd8970
R10: 0000000000000000 R11: 0000000000000202 R12: 00007ffca0df9308
R13: 0000000000000001 R14: 0000000000000000 R15: 0000000000000003
---[ end trace 2f501b01b348f114 ]---
ACPI: Preparing to enter system sleep state S5
reboot: Power down

Changes since v1:
```

```
      - Rebase to latest drm-tip

    Fixes: 0cfab4cb3c4e ("drm/i915/pxp: Enable PXP power management")
    Suggested-by: Lee Shawn C <shawn.c.lee@intel.com>
    Signed-off-by: Juston Li <juston.li@intel.com>
    Reviewed-by: Daniele Ceraolo Spurio <daniele.ceraolospurio@intel.com>
    Signed-off-by: John Harrison <John.C.Harrison@Intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220106200236.489656-2-
juston.li@intel.com
    (cherry picked from commit 57ded5fc98b11d76dae505ca3591b61c9dbbbda7)
    Signed-off-by: Tvrtko Ursulin <tvrtko.ursulin@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8f1027cb135276289a6141c9635014a615bdf981
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Thu Jan 6 10:46:11 2022 +0100

    x86/hyperv: Properly deal with empty cpumasks in hyperv_flush_tlb_multi()

    [ Upstream commit 51500b71d500f251037ed339047a4d9e7d7e295b ]

    KASAN detected the following issue:

     BUG: KASAN: slab-out-of-bounds in hyperv_flush_tlb_multi+0xf88/0x1060
     Read of size 4 at addr ffff8880011ccbc0 by task kcompactd0/33

     CPU: 1 PID: 33 Comm: kcompactd0 Not tainted 5.14.0-39.el9.x86_64+debug #1
     Hardware name: Microsoft Corporation Virtual Machine/Virtual Machine,
         BIOS Hyper-V UEFI Release v4.0 12/17/2019
     Call Trace:
      dump_stack_lvl+0x57/0x7d
      print_address_description.constprop.0+0x1f/0x140
      ? hyperv_flush_tlb_multi+0xf88/0x1060
      __kasan_report.cold+0x7f/0x11e
      ? hyperv_flush_tlb_multi+0xf88/0x1060
      kasan_report+0x38/0x50
      hyperv_flush_tlb_multi+0xf88/0x1060
      flush_tlb_mm_range+0x1b1/0x200
      ptep_clear_flush+0x10e/0x150
    ...
     Allocated by task 0:
      kasan_save_stack+0x1b/0x40
      __kasan_kmalloc+0x7c/0x90
      hv_common_init+0xae/0x115
      hyperv_init+0x97/0x501
      apic_intr_mode_init+0xb3/0x1e0
      x86_late_time_init+0x92/0xa2
      start_kernel+0x338/0x3eb
      secondary_startup_64_no_verify+0xc2/0xcb

     The buggy address belongs to the object at ffff8880011cc800
      which belongs to the cache kmalloc-1k of size 1024
     The buggy address is located 960 bytes inside of
      1024-byte region [ffff8880011cc800, ffff8880011ccc00)

    'hyperv_flush_tlb_multi+0xf88/0x1060' points to
    hv_cpu_number_to_vp_number() and '960 bytes' means we're trying to get
    VP_INDEX for CPU#240. 'nr_cpus' here is exactly 240 so we're trying to
    access past hv_vp_index's last element. This can (and will) happen
    when 'cpus' mask is empty and cpumask_last() will return '>=nr_cpus'.

    Commit ad0a6bad4475 ("x86/hyperv: check cpu mask after interrupt has
    been disabled") tried to deal with empty cpumask situation but
    apparently didn't fully fix the issue.
```

```
    'cpus' cpumask which is passed to hyperv_flush_tlb_multi() is
    'mm_cpumask(mm)' (which is '&mm->cpu_bitmap'). This mask changes every
    time the particular mm is scheduled/unscheduled on some CPU (see
    switch_mm_irqs_off()), disabling IRQs on the CPU which is performing remote
    TLB flush has zero influence on whether the particular process can get
    scheduled/unscheduled on _other_ CPUs so e.g. in the case where the mm was
    scheduled on one other CPU and got unscheduled during
    hyperv_flush_tlb_multi()'s execution will lead to cpumask becoming empty.

    It doesn't seem that there's a good way to protect 'mm_cpumask(mm)'
    from changing during hyperv_flush_tlb_multi()'s execution. It would be
    possible to copy it in the very beginning of the function but this is a
    waste. It seems we can deal with changing cpumask just fine.

    When 'cpus' cpumask changes during hyperv_flush_tlb_multi()'s
    execution, there are two possible issues:
    - 'Under-flushing': we will not flush TLB on a CPU which got added to
    the mask while hyperv_flush_tlb_multi() was already running. This is
    not a problem as this is equal to mm getting scheduled on that CPU
    right after TLB flush.
    - 'Over-flushing': we may flush TLB on a CPU which is already cleared
    from the mask. First, extra TLB flush preserves correctness. Second,
    Hyper-V's TLB flush hypercall takes 'mm->pgd' argument so Hyper-V may
    avoid the flush if CR3 doesn't match.

    Fix the immediate issue with cpumask_last()/hv_cpu_number_to_vp_number()
    and remove the pointless cpumask_empty() check from the beginning of the
    function as it really doesn't protect anything. Also, avoid the hypercall
    altogether when 'flush->processor_mask' ends up being empty.

    Fixes: ad0a6bad4475 ("x86/hyperv: check cpu mask after interrupt has been disabled")
    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Reviewed-by: Michael Kelley <mikelley@microsoft.com>
    Link: https://lore.kernel.org/r/20220106094611.1404218-1-vkuznets@redhat.com
    Signed-off-by: Wei Liu <wei.liu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c7aaaa5fa9f21d101825c68b6c09f2d0aecd371
Author: J. Bruce Fields <bfields@redhat.com>
Date:   Wed Jan 5 14:15:03 2022 -0500

    nfsd: fix crash on COPY_NOTIFY with special stateid

    [ Upstream commit 074b07d94e0bb6ddce5690a9b7e2373088e8b33a ]

    RTM says "If the special ONE stateid is passed to
    nfs4_preprocess_stateid_op(), it returns status=0 but does not set
    *cstid. nfsd4_copy_notify() depends on stid being set if status=0, and
    thus can crash if the client sends the right COPY_NOTIFY RPC."

    RFC 7862 says "The cna_src_stateid MUST refer to either open or locking
    states provided earlier by the server.  If it is invalid, then the
    operation MUST fail."

    The RFC doesn't specify an error, and the choice doesn't matter much as
    this is clearly illegal client behavior, but bad_stateid seems
    reasonable.

    Simplest is just to guarantee that nfs4_preprocess_stateid_op, called
    with non-NULL cstid, errors out if it can't return a stateid.

    Reported-by: rtm@csail.mit.edu
    Fixes: 624322f1adc5 ("NFSD add COPY_NOTIFY operation")
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
```

```
    Reviewed-by: Olga Kornievskaia <kolga@netapp.com>
    Tested-by: Olga Kornievskaia <kolga@netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a222f2c34b237a75d634aec9b0ef8a2af9fbe22
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Fri Dec 24 14:22:28 2021 -0500

    Revert "nfsd: skip some unnecessary stats in the v4 case"

    [ Upstream commit 58f258f65267542959487dbe8b5641754411843d ]

    On the wire, I observed NFSv4 OPEN(CREATE) operations sometimes
    returning a reasonable-looking value in the cinfo.before field and
    zero in the cinfo.after field.

    RFC 8881 Section 10.8.1 says:
    > When a client is making changes to a given directory, it needs to
    > determine whether there have been changes made to the directory by
    > other clients.  It does this by using the change attribute as
    > reported before and after the directory operation in the associated
    > change_info4 value returned for the operation.

    and

    > ... The post-operation change
    > value needs to be saved as the basis for future change_info4
    > comparisons.

    A good quality client implementation therefore saves the zero
    cinfo.after value. During a subsequent OPEN operation, it will
    receive a different non-zero value in the cinfo.before field for
    that directory, and it will incorrectly believe the directory has
    changed, triggering an undesirable directory cache invalidation.

    There are filesystem types where fs_supports_change_attribute()
    returns false, tmpfs being one. On NFSv4 mounts, this means the
    fh_getattr() call site in fill_pre_wcc() and fill_post_wcc() is
    never invoked. Subsequently, nfsd4_change_attribute() is invoked
    with an uninitialized @stat argument.

    In fill_pre_wcc(), @stat contains stale stack garbage, which is
    then placed on the wire. In fill_post_wcc(), ->fh_post_wc is all
    zeroes, so zero is placed on the wire. Both of these values are
    meaningless.

    This fix can be applied immediately to stable kernels. Once there
    are more regression tests in this area, this optimization can be
    attempted again.

    Fixes: 428a23d2bf0c ("nfsd: skip some unnecessary stats in the v4 case")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fbd1523f542b3269805fe7b37e8a8ebd33607b2a
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Tue Dec 28 12:35:43 2021 -0500

    NFSD: Fix verifier returned in stable WRITEs

    [ Upstream commit f11ad7aa653130b71e2e89bed207f387718216d5 ]

    RFC 8881 explains the purpose of the write verifier this way:

    > The final portion of the result is the field writeverf. This field
```

```
    > is the write verifier and is a cookie that the client can use to
    > determine whether a server has changed instance state (e.g., server
    > restart) between a call to WRITE and a subsequent call to either
    > WRITE or COMMIT.

    But then it says:

    > This cookie MUST be unchanged during a single instance of the
    > NFSv4.1 server and MUST be unique between instances of the NFSv4.1
    > server. If the cookie changes, then the client MUST assume that
    > any data written with an UNSTABLE4 value for committed and an old
    > writeverf in the reply has been lost and will need to be
    > recovered.

    RFC 1813 has similar language for NFSv3. NFSv2 does not have a write
    verifier since it doesn't implement the COMMIT procedure.

    Since commit 19e0663ff9bc ("nfsd: Ensure sampling of the write
    verifier is atomic with the write"), the Linux NFS server has
    returned a boot-time-based verifier for UNSTABLE WRITEs, but a zero
    verifier for FILE_SYNC and DATA_SYNC WRITEs. FILE_SYNC and DATA_SYNC
    WRITEs are not followed up with a COMMIT, so there's no need for
    clients to compare verifiers for stable writes.

    However, by returning a different verifier for stable and unstable
    writes, the above commit puts the Linux NFS server a step farther
    out of compliance with the first MUST above. At least one NFS client
    (FreeBSD) noticed the difference, making this a potential
    regression.

    Reported-by: Rick Macklem <rmacklem@uoguelph.ca>
    Link: https://lore.kernel.org/linux-
nfs/YQXPR0101MB096857EEACF04A6DF1FC6D9BDD749@YQXPR0101MB0968.CANPRD01.PROD.OUTLOOK.COM/T/
    Fixes: 19e0663ff9bc ("nfsd: Ensure sampling of the write verifier is atomic with the
write")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9547539fa15054011f95c8bae7acd13c12b71f8
Author: Vincent Chen <vincent.chen@sifive.com>
Date:   Mon Dec 27 11:05:14 2021 +0800

    KVM: RISC-V: Avoid spurious virtual interrupts after clearing hideleg CSR

    [ Upstream commit 33e5b5746cc2336660c8710ba109d9a3923627b5 ]

    When the last VM is terminated, the host kernel will invoke function
    hardware_disable_nolock() on each CPU to disable the related virtualization
    functions. Here, RISC-V currently only clears hideleg CSR and hedeleg CSR.
    This behavior will cause the host kernel to receive spurious interrupts if
    hvip CSR has pending interrupts and the corresponding enable bits in vsie
    CSR are asserted. To avoid it, hvip CSR and vsie CSR must be cleared
    before clearing hideleg CSR.

    Fixes: 99cdc6c18c2d ("RISC-V: Add initial skeletal KVM support")
    Signed-off-by: Vincent Chen <vincent.chen@sifive.com>
    Reviewed-by: Anup Patel <anup.patel@wdc.com>
    Signed-off-by: Anup Patel <anup.patel@wdc.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce16d4b7e5f6076c2a5ab52e3826201ed20d9ac9
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:46:05 2021 +0100

    PCI: mvebu: Fix support for DEVCAP2, DEVCTL2 and LNKCTL2 registers on emulated bridge
```

```
    [ Upstream commit 4ab34548c55fbbb3898306a47dfaccd4860e1ccb ]

    Armada XP and new hardware supports access to DEVCAP2, DEVCTL2 and LNKCTL2
    configuration registers of PCIe core via PCIE_CAP_PCIEXP. So export them
    via emulated software root bridge.

    Pre-XP hardware does not support these registers and returns zeros.

    Link: https://lore.kernel.org/r/20211125124605.25915-16-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 004408c5b7b47f80afec75ade3e788b65c1c8b99
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:46:04 2021 +0100

    PCI: mvebu: Fix support for PCI_EXP_RTSTA on emulated bridge

    [ Upstream commit 838ff44a398ff47fe9b924961d91aee325821220 ]

    PME Status bit in Root Status Register (PCIE_RC_RTSTA_OFF) is read-only and
    can be cleared only by writing 0b to the Interrupt Cause RW0C register
    (PCIE_INT_CAUSE_OFF).

    Link: https://lore.kernel.org/r/20211125124605.25915-15-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9dd0d0efecee653187536f26c902f2bc21ecf06
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:46:03 2021 +0100

    PCI: mvebu: Fix support for PCI_EXP_DEVCTL on emulated bridge

    [ Upstream commit ecae073e393e65ee7be7ebf3fdd5258ab99f1636 ]

    Comment in Armada 370 functional specification is misleading.
    PCI_EXP_DEVCTL_*RE bits are supported and configures receiving of error
    interrupts.

    Link: https://lore.kernel.org/r/20211125124605.25915-14-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 802d9ee9cbd3cb0430ea72de2a836a861e2710c4
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:46:02 2021 +0100

    PCI: mvebu: Fix support for PCI_BRIDGE_CTL_BUS_RESET on emulated bridge

    [ Upstream commit d75404cc08832206f173668bd35391c581fea121 ]

    Hardware supports PCIe Hot Reset via PCIE_CTRL_OFF register. Use it for
    implementing PCI_BRIDGE_CTL_BUS_RESET bit of PCI_BRIDGE_CONTROL register on
    emulated bridge.

    With this change the function pci_reset_secondary_bus() starts working and
    can reset connected PCIe card.
```

```
    Link: https://lore.kernel.org/r/20211125124605.25915-13-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4523e727c349f22e09086e5af3c91f3ec46fe59e
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:45:59 2021 +0100

    PCI: mvebu: Setup PCIe controller to Root Complex mode

    [ Upstream commit df08ac016124bd88b8598ac0599d7b89c0642774 ]

    This driver operates only in Root Complex mode, so ensure that hardware is
    properly configured in Root Complex mode.

    Link: https://lore.kernel.org/r/20211125124605.25915-10-pali@kernel.org
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7cde9bf0731688896831f90da9fe755f44a6d5e0
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:46:01 2021 +0100

    PCI: mvebu: Fix configuring secondary bus of PCIe Root Port via emulated bridge

    [ Upstream commit 91a8d79fc797d3486ae978beebdfc55261c7d65b ]

    It looks like that mvebu PCIe controller has for each PCIe link fully
    independent PCIe host bridge and so every PCIe Root Port is isolated not
    only on its own bus but also isolated from each others. But in past device
    tree structure was defined to put all PCIe Root Ports (as PCI Bridge
    devices) into one root bus 0 and this bus is emulated by pci-mvebu.c
    driver.

    Probably reason for this decision was incorrect understanding of PCIe
    topology of these Armada SoCs and also reason of misunderstanding how is
    PCIe controller generating Type 0 and Type 1 config requests (it is fully
    different compared to other drivers). Probably incorrect setup leaded to
    very surprised things like having PCIe Root Port (PCI Bridge device, with
    even incorrect Device Class set to Memory Controller) and the PCIe device
    behind the Root Port on the same PCI bus, which obviously was needed to
    somehow hack (as these two devices cannot be in reality on the same bus).

    Properly set mvebu local bus number and mvebu local device number based on
    PCI Bridge secondary bus number configuration. Also correctly report
    configured secondary bus number in config space. And explain in driver
    comment why this setup is correct.

    Link: https://lore.kernel.org/r/20211125124605.25915-12-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3de91c80b70ab84bc298930199c1a0a5b607fbf3
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:45:56 2021 +0100

    PCI: mvebu: Fix support for bus mastering and PCI_COMMAND on emulated bridge

    [ Upstream commit e42b85583719adb87ab88dc7bcd41b38011f7d11 ]
```

```
    According to PCI specifications bits [0:2] of Command Register, this should
    be by default disabled on reset. So explicitly disable these bits at early
    beginning of driver initialization.

    Also remove code which unconditionally enables all 3 bits and let kernel
    code (via pci_set_master() function) to handle bus mastering of PCI Bridge
    via emulated PCI_COMMAND on emulated bridge.

    Adjust existing functions mvebu_pcie_handle_iobase_change() and
    mvebu_pcie_handle_membase_change() to handle PCI_IO_BASE and PCI_MEM_BASE
    registers correctly even when bus mastering on emulated bridge is disabled.

    Link: https://lore.kernel.org/r/20211125124605.25915-7-pali@kernel.org
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d9bfeaab65b30f29f32c3b40859b2c267af792e4
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:45:57 2021 +0100

    PCI: mvebu: Do not modify PCI IO type bits in conf_write

    [ Upstream commit 2cf150216e5b5619d7c25180ccf2cc8ac7bebc13 ]

    PCI IO type bits are already initialized in mvebu_pci_bridge_emul_init()
    function and only when IO support is enabled. These type bits are read-only
    and pci-bridge-emul.c code already does not allow to modify them from upper
    layers.

    When IO support is disabled then all IO registers should be read-only and
    return zeros. Therefore do not modify PCI IO type bits in
    mvebu_pci_bridge_emul_base_conf_write() callback.

    Link: https://lore.kernel.org/r/20211125124605.25915-8-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e7e52bc070216da620086fbe0a3d89745fa9ad04
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 13:45:52 2021 +0100

    PCI: mvebu: Check for errors from pci_bridge_emul_init() call

    [ Upstream commit 5d18d702e5c9309f4195653475c7a7fdde4ca71f ]

    Function pci_bridge_emul_init() may fail so correctly check for errors.

    Link: https://lore.kernel.org/r/20211125124605.25915-3-pali@kernel.org
    Fixes: 1f08673eef12 ("PCI: mvebu: Convert to PCI emulated bridge config space")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cace5de27d97e5d02ab31d1e0dec9388a0fc4854
Author: Dario Binacchi <dariobin@libero.it>
Date:   Sun Dec 12 21:14:48 2021 -0800

    Input: ti_am335x_tsc - fix STEPCONFIG setup for Z2

    [ Upstream commit 6bfeb6c21e1bdc11c328b7d996d20f0f73c6b9b0 ]
```

```
    The Z2 step configuration doesn't erase the SEL_INP_SWC_3_0 bit-field
    before setting the ADC channel. This way its value could be corrupted by
    the ADC channel selected for the Z1 coordinate.

    Fixes: 8c896308feae ("input: ti_am335x_adc: use only FIFO0 and clean up a little")
    Signed-off-by: Dario Binacchi <dariobin@libero.it>
    Link: https://lore.kernel.org/r/20211212125358.14416-3-dariobin@libero.it
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 757609438d0b9b90578adb4b39eda4ac788e3e20
Author: Dario Binacchi <dariobin@libero.it>
Date:   Sun Dec 12 21:14:35 2021 -0800

    Input: ti_am335x_tsc - set ADCREFM for X configuration

    [ Upstream commit 73cca71a903202cddc8279fc76b2da4995da5bea ]

    As reported by the STEPCONFIG[1-16] registered field descriptions of the
    TI reference manual, for the ADC "in single ended, SEL_INM_SWC_3_0 must
    be 1xxx".

    Unlike the Y and Z coordinates, this bit has not been set for the step
    configuration registers used to sample the X coordinate.

    Fixes: 1b8be32e6914 ("Input: add support for TI Touchscreen controller")
    Signed-off-by: Dario Binacchi <dariobin@libero.it>
    Link: https://lore.kernel.org/r/20211212125358.14416-2-dariobin@libero.it
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fda2c6406924e1f035cf0676faedbb01dec710e4
Author: Beau Belgrave <beaub@linux.microsoft.com>
Date:   Thu Sep 30 15:38:21 2021 -0700

    tracing: Do not let synth_events block other dyn_event systems during create

    [ Upstream commit 4f67cca70c0f615e9cfe6ac42244f3416ec60877 ]

    synth_events is returning -EINVAL if the dyn_event create command does
    not contain ' \t'. This prevents other systems from getting called back.
    synth_events needs to return -ECANCELED in these cases when the command
    is not targeting the synth_event system.

    Link: https://lore.kernel.org/linux-trace-devel/20210930223821.11025-1-
beaub@linux.microsoft.com

    Fixes: c9e759b1e8456 ("tracing: Rework synthetic event command parsing")
    Reviewed-by: Masami Hiramatsu <mhiramat@kernel.org>
    Signed-off-by: Beau Belgrave <beaub@linux.microsoft.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e80e3bc9ae346ed93dc205dcacb1e756dc4fb41
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Wed Nov 17 23:05:23 2021 +0100

    i3c/master/mipi-i3c-hci: Fix a potentially infinite loop in 'hci_dat_v1_get_index()'

    [ Upstream commit 3f43926f271287fb1744c9ac9ae1122497f2b0c2 ]

    The code in 'hci_dat_v1_get_index()' really looks like a hand coded version
    of 'for_each_set_bit()', except that a +1 is missing when searching for the
    next set bit.
```

```
    This really looks odd and it seems that it will loop until 'dat_w0_read()'
    returns the expected result.

    So use 'for_each_set_bit()' instead. It is less verbose and should be more
    correct.

    Fixes: 9ad9a52cce28 ("i3c/master: introduce the mipi-i3c-hci driver")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Acked-by: Nicolas Pitre <npitre@baylibre.com>
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link:
https://lore.kernel.org/r/0cdf3cb10293ead1acd271fdb8a70369c298c082.1637186628.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2635263a750b7760cbec170a5e1c928e406378f
Author: Jamie Iles <quic_jiles@quicinc.com>
Date:   Wed Sep 22 17:56:00 2021 +0100

    i3c: fix incorrect address slot lookup on 64-bit

    [ Upstream commit f18f98110f2b179792cb70d85cba697320a3790f ]

    The address slot bitmap is an array of unsigned long's which are the
    same size as an int on 32-bit platforms but not 64-bit.  Loading the
    bitmap into an int could result in the incorrect status being returned
    for a slot and slots being reported as the wrong status.

    Fixes: 3a379bbcea0a ("i3c: Add core I3C infrastructure")
    Cc: Boris Brezillon <bbrezillon@kernel.org>
    Cc: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Jamie Iles <quic_jiles@quicinc.com>
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/20210922165600.179394-1-quic_jiles@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 009fb2aba6173a9264f82b1e50610d14fbd7c249
Author: Hou Wenlong <houwenlong93@linux.alibaba.com>
Date:   Tue Nov 2 17:15:32 2021 +0800

    KVM: x86: Exit to userspace if emulation prepared a completion callback

    [ Upstream commit adbfb12d4c4517a8adde23a7fc46538953d56eea ]

    em_rdmsr() and em_wrmsr() return X86EMUL_IO_NEEDED if MSR accesses
    required an exit to userspace. However, x86_emulate_insn() doesn't return
    X86EMUL_*, so x86_emulate_instruction() doesn't directly act on
    X86EMUL_IO_NEEDED; instead, it looks for other signals to differentiate
    between PIO, MMIO, etc. causing RDMSR/WRMSR emulation not to
    exit to userspace now.

    Nevertheless, if the userspace_msr_exit_test testcase in selftests
    is changed to test RDMSR/WRMSR with a forced emulation prefix,
    the test passes.  What happens is that first userspace exit
    information is filled but the userspace exit does not happen.
    Because x86_emulate_instruction() returns 1, the guest retries
    the instruction---but this time RIP has already been adjusted
    past the forced emulation prefix, so the guest executes RDMSR/WRMSR
    and the userspace exit finally happens.

    Since the X86EMUL_IO_NEEDED path has provided a complete_userspace_io
    callback, x86_emulate_instruction() can just return 0 if the
    callback is not NULL. Then RDMSR/WRMSR instruction emulation will
    exit to userspace directly, without the RDMSR/WRMSR vmexit.
```

```
    Fixes: 1ae099540e8c7 ("KVM: x86: Allow deflecting unknown MSR accesses to user space")
    Signed-off-by: Hou Wenlong <houwenlong93@linux.alibaba.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Message-Id:
<56f9df2ee5c05a81155e2be366c9dc1f7adc8817.1635842679.git.houwenlong93@linux.alibaba.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3fdfb82a3b98f4dacc9c6be3be5d904fa7b9c7ac
Author: Sean Christopherson <seanjc@google.com>
Date:   Tue Nov 2 17:15:29 2021 +0800

    KVM: x86: Handle 32-bit wrap of EIP for EMULTYPE_SKIP with flat code seg

    [ Upstream commit 5e854864ee4384736f27a986633bae21731a4e4e ]

    Truncate the new EIP to a 32-bit value when handling EMULTYPE_SKIP as the
    decode phase does not truncate _eip.  Wrapping the 32-bit boundary is
    legal if and only if CS is a flat code segment, but that check is
    implicitly handled in the form of limit checks in the decode phase.

    Opportunstically prepare for a future fix by storing the result of any
    truncation in "eip" instead of "_eip".

    Fixes: 1957aa63be53 ("KVM: VMX: Handle single-step #DB for EMULTYPE_SKIP on EPT
misconfig")
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Message-Id:
<093eabb1eab2965201c9b018373baf26ff256d85.1635842679.git.houwenlong93@linux.alibaba.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a538bd5652394493df6ff59110ecec00e9245f8
Author: Lai Jiangshan <laijs@linux.alibaba.com>
Date:   Mon Nov 8 20:43:53 2021 +0800

    KVM: X86: Ensure that dirty PDPTRs are loaded

    [ Upstream commit 2c5653caecc4807b8abfe9c41880ac38417be7bf ]

    For VMX with EPT, dirty PDPTRs need to be loaded before the next vmentry
    via vmx_load_mmu_pgd()

    But not all paths that call load_pdptrs() will cause vmx_load_mmu_pgd()
    to be invoked.  Normally, kvm_mmu_reset_context() is used to cause
    KVM_REQ_LOAD_MMU_PGD, but sometimes it is skipped:

    * commit d81135a57aa6("KVM: x86: do not reset mmu if CR0.CD and
    CR0.NW are changed") skips kvm_mmu_reset_context() after load_pdptrs()
    when changing CR0.CD and CR0.NW.

    * commit 21823fbda552("KVM: x86: Invalidate all PGDs for the current
    PCID on MOV CR3 w/ flush") skips KVM_REQ_LOAD_MMU_PGD after
    load_pdptrs() when rewriting the CR3 with the same value.

    * commit a91a7c709600("KVM: X86: Don't reset mmu context when
    toggling X86_CR4_PGE") skips kvm_mmu_reset_context() after
    load_pdptrs() when changing CR4.PGE.

    Fixes: d81135a57aa6 ("KVM: x86: do not reset mmu if CR0.CD and CR0.NW are changed")
    Fixes: 21823fbda552 ("KVM: x86: Invalidate all PGDs for the current PCID on MOV CR3 w/
flush")
    Fixes: a91a7c709600 ("KVM: X86: Don't reset mmu context when toggling X86_CR4_PGE")
    Signed-off-by: Lai Jiangshan <laijs@linux.alibaba.com>
    Message-Id: <20211108124407.12187-2-jiangshanlai@gmail.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
```

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

commit fffbed22272700bf06a1cda447b06044e1fb941f
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Oct 8 19:12:19 2021 -0700

    KVM: VMX: Read Posted Interrupt "control" exactly once per loop iteration

    [ Upstream commit cfb0e1306a3790eb055ebf7cdb7b0ee8a23e9b6e ]

    Use READ_ONCE() when loading the posted interrupt descriptor control
    field to ensure "old" and "new" have the same base value.  If the
    compiler emits separate loads, and loads into "new" before "old", KVM
    could theoretically drop the ON bit if it were set between the loads.

    Fixes: 28b835d60fcc ("KVM: Update Posted-Interrupts Descriptor when vCPU is preempted")
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20211009021236.4122790-27-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69bbdf32f04178cef945945439181296dd51d23a
Author: Sean Christopherson <seanjc@google.com>
Date:   Fri Oct 8 19:11:56 2021 -0700

    KVM: s390: Ensure kvm_arch_no_poll() is read once when blocking vCPU

    [ Upstream commit 6f390916c4fb359507d9ac4bf1b28a4f8abee5c0 ]

    Wrap s390's halt_poll_max_steal with READ_ONCE and snapshot the result of
    kvm_arch_no_poll() in kvm_vcpu_block() to avoid a mostly-theoretical,
    largely benign bug on s390 where the result of kvm_arch_no_poll() could
    change due to userspace modifying halt_poll_max_steal while the vCPU is
    blocking.  The bug is largely benign as it will either cause KVM to skip
    updating halt-polling times (no_poll toggles false=>true) or to update
    halt-polling times with a slightly flawed block_ns.

    Note, READ_ONCE is unnecessary in the current code, add it in case the
    arch hook is ever inlined, and to provide a hint that userspace can
    change the param at will.

    Fixes: 8b905d28ee17 ("KVM: s390: provide kvm_arch_no_poll function")
    Reviewed-by: Christian Borntraeger <borntraeger@de.ibm.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20211009021236.4122790-4-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa59606d5b93cf0c4c4b86a40fffd32f4453c6ed
Author: Paolo Bonzini <pbonzini@redhat.com>
Date:   Tue Nov 16 09:32:47 2021 -0500

    KVM: VMX: Don't unblock vCPU w/ Posted IRQ if IRQs are disabled in guest

    [ Upstream commit 1831fa44df743a7cdffdf1c12c799bf6f3c12b8c ]

    Don't configure the wakeup handler when a vCPU is blocking with IRQs
    disabled, in which case any IRQ, posted or otherwise, should not be
    recognized and thus should not wake the vCPU.

    Fixes: bf9f6ac8d749 ("KVM: Update Posted-Interrupts Descriptor when vCPU is blocked")
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20211009021236.4122790-2-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 64028e0f7ce43c3b2f6fdfe99283fc6063776c23
Author: Hector Martin <marcan@marcan.st>
Date:   Wed Nov 17 23:00:44 2021 +0900

    PCI: apple: Fix REFCLK1 enable/poll logic

    [ Upstream commit 75d36df6807838f3c826c21c0fa51cdc079667d1 ]

    REFCLK1 has req/ack bits that need to be programmed, just like REFCLK0.

    Link: https://lore.kernel.org/r/20211117140044.193865-1-marcan@marcan.st
    Fixes: 1e33888fbe44 ("PCI: apple: Add initial hardware bring-up")
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Acked-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2bda5a25e6022e89008de44789bafc01ce69b8a
Author: Pali Rohár <pali@kernel.org>
Date:   Thu Nov 25 17:01:47 2021 +0100

    PCI: aardvark: Fix checking for MEM resource type

    [ Upstream commit 2070b2ddea89f5b604fac3d27ade5cb6d19a5706 ]

    IORESOURCE_MEM_64 is not a resource type but a type flag.

    Remove incorrect check for type IORESOURCE_MEM_64.

    Link: https://lore.kernel.org/r/20211125160148.26029-2-kabel@kernel.org
    Fixes: 64f160e19e92 ("PCI: aardvark: Configure PCIe resources from 'ranges' DT
property")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f660c32729f78d3a24cd0ef026835c10d8d7c5fa
Author: Tim Harvey <tharvey@gateworks.com>
Date:   Mon Nov 1 11:02:43 2021 -0700

    PCI: dwc: Do not remap invalid res

    [ Upstream commit 6e5ebc96ec651b67131f816d7e3bf286c635e749 ]

    On imx6 and perhaps others when pcie probes you get a:
    imx6q-pcie 33800000.pcie: invalid resource

    This occurs because the atu is not specified in the DT and as such it
    should not be remapped.

    Link: https://lore.kernel.org/r/20211101180243.23761-1-tharvey@gateworks.com
    Fixes: 281f1f99cf3a ("PCI: dwc: Detect number of iATU windows")
    Signed-off-by: Tim Harvey <tharvey@gateworks.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Rob Herring <robh@kernel.org>
    Acked-by: Richard Zhu <hongxing.zhu@nxp.com>
    Cc: Richard Zhu <hongxing.zhu@nxp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 599735b05d411e1dbad86c3e71fa631f5c4ebaa6
Author: Marek Vasut <marek.vasut+renesas@gmail.com>
Date:   Mon Nov 15 21:46:41 2021 +0100
```

```
    PCI: rcar: Check if device is runtime suspended instead of __clk_is_enabled()

    [ Upstream commit d2a14b54989e9ccea8401895fdfbc213bd1f56af ]

    Replace __clk_is_enabled() with pm_runtime_suspended(),
    as __clk_is_enabled() was checking the wrong bus clock
    and caused the following build error too:
      arm-linux-gnueabi-ld: drivers/pci/controller/pcie-rcar-host.o: in function
`rcar_pcie_aarch32_abort_handler':
      pcie-rcar-host.c:(.text+0xdd0): undefined reference to `__clk_is_enabled'

    Link: https://lore.kernel.org/r/20211115204641.12941-1-marek.vasut@gmail.com
    Fixes: a115b1bd3af0 ("PCI: rcar: Add L1 link state fix into data abort hook")
    Signed-off-by: Marek Vasut <marek.vasut+renesas@gmail.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Acked-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Arnd Bergmann <arnd@arndb.de>
    Cc: Bjorn Helgaas <bhelgaas@google.com>
    Cc: Geert Uytterhoeven <geert+renesas@glider.be>
    Cc: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Cc: Stephen Boyd <sboyd@kernel.org>
    Cc: Wolfram Sang <wsa@the-dreams.de>
    Cc: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Cc: linux-renesas-soc@vger.kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 26dc73017ec7626d3660f4761c41a0c2d531c609
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sat Nov 6 18:44:52 2021 +0100

    PCI: qcom: Fix an error handling path in 'qcom_pcie_probe()'

    [ Upstream commit 4e0e90539bb0e6c0ca3768c642df9eed2118a8bb ]

    If 'of_device_get_match_data()' fails, previous 'pm_runtime_get_sync()/
    pm_runtime_enable()' should be undone.

    To fix it, the easiest is to move this block of code before the memory
    allocations and the pm_runtime_xxx calls.

    Link:
https://lore.kernel.org/r/4d03c636193f64907c8dacb17fa71ed05fd5f60c.1636220582.git.christophe
.jaillet@wanadoo.fr
    Fixes: b89ff410253d ("PCI: qcom: Replace ops with struct pcie_cfg in pcie match data")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3b314b10526218f89b0fe003168355ec2343f33
Author: Jianjun Wang <jianjun.wang@mediatek.com>
Date:   Fri Oct 15 14:36:02 2021 +0800

    PCI: mediatek-gen3: Disable DVFSRC voltage request

    [ Upstream commit ab344fd43f2958726d17d651c0cb692c67dca382 ]

    When the DVFSRC (dynamic voltage and frequency scaling resource collector)
    feature is not implemented, the PCIe hardware will assert a voltage request
    signal when exit from the L1 PM Substates to request a specific Vcore
    voltage, but cannot receive the voltage ready signal, which will cause
    the link to fail to exit the L1 PM Substates.

    Disable DVFSRC voltage request by default, we need to find a common way to
```

```
    enable it in the future.

    Link: https://lore.kernel.org/r/20211015063602.29058-1-jianjun.wang@mediatek.com
    Fixes: d3bf75b579b9 ("PCI: mediatek-gen3: Add MediaTek Gen3 driver for MT8192")
    Tested-by: Qizhong Cheng <qizhong.cheng@mediatek.com>
    Signed-off-by: Jianjun Wang <jianjun.wang@mediatek.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Tzung-Bi Shih <tzungbi@google.com>
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7acc9d00d0f315651120f4f3ebaf27610a8c6700
Author: Eric W. Biederman <ebiederm@xmission.com>
Date:   Mon Nov 15 11:55:57 2021 -0600

    signal: In get_signal test for signal_group_exit every time through the loop

    [ Upstream commit e7f7c99ba911f56bc338845c1cd72954ba591707 ]

    Recently while investigating a problem with rr and signals I noticed
    that siglock is dropped in ptrace_signal and get_signal does not jump
    to relock.

    Looking farther to see if the problem is anywhere else I see that
    do_signal_stop also returns if signal_group_exit is true.  I believe
    that test can now never be true, but it is a bit hard to trace
    through and be certain.

    Testing signal_group_exit is not expensive, so move the test for
    signal_group_exit into the for loop inside of get_signal to ensure
    the test is never skipped improperly.

    This has been a potential problem since I added the test for
    signal_group_exit was added.

    Fixes: 35634ffa1751 ("signal: Always notice exiting tasks")
    Reviewed-by: Kees Cook <keescook@chromium.org>
    Link: https://lkml.kernel.org/r/875yssekcd.fsf_-_@email.froward.int.ebiederm.org
    Signed-off-by: "Eric W. Biederman" <ebiederm@xmission.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3a2e59a4e9af9f153bffe08597d589d628ccfe3
Author: Conor Dooley <conor.dooley@microchip.com>
Date:   Fri Dec 17 09:33:12 2021 +0000

    mailbox: change mailbox-mpfs compatible string

    [ Upstream commit f10b1fc0161cd99e54c5687fcc63368aa255e05e ]

    The Polarfire SoC is currently using two different compatible string
    prefixes. Fix this by changing "polarfire-soc-*" strings to "mpfs-*" in
    its system controller in order to match the compatible string used in
    the soc binding and device tree.

    Signed-off-by: Conor Dooley <conor.dooley@microchip.com>
    Reviewed-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1c18c397cd9af4f281108a635d579efb3f1f670
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Dec 24 08:21:03 2021 +0000

    phy: mediatek: Fix missing check in mtk_mipi_tx_probe
```

```
    [ Upstream commit 399c91c3f30531593e5ff6ca7b53f47092128669 ]

    The of_device_get_match_data() function may return NULL.
    Add check to prevent potential null dereference.

    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20211224082103.7658-1-linmq006@gmail.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8423ec05aff2b41ca657e56fb3dd7ca0da566cb1
Author: Ohad Sharabi <osharabi@habana.ai>
Date:   Mon Nov 22 12:23:51 2021 +0200

    habanalabs: skip read fw errors if dynamic descriptor invalid

    [ Upstream commit 4fac990f604e6c10538026835a8a30f3c1b6fcf5 ]

    Reporting FW errors involves reading of the error registers.

    In case we have a corrupted FW descriptor we cannot do that since the
    dynamic scratchpad is potentially corrupted as well and may cause kernel
    crush when attempting access to a corrupted register offset.

    Signed-off-by: Ohad Sharabi <osharabi@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1bc4b7489c7f5a7ac92eae051d4cb5fa37f1654
Author: Dani Liberman <dliberman@habana.ai>
Date:   Thu Oct 14 22:38:41 2021 +0300

    habanalabs: change wait for interrupt timeout to 64 bit

    [ Upstream commit 48f31169830f589e4c7ac475ccc7414951ded3f0 ]

    In order to increase maximum wait-for-interrupt timeout, change it
    to 64 bit variable. This wait is used only by newer ASICs, so no
    problem in changing this interface at this time.

    Signed-off-by: Dani Liberman <dliberman@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 160fc1643e2d7347b66694ebe3aab2010517d5cf
Author: Tzung-Bi Shih <tzungbi@google.com>
Date:   Fri Dec 24 14:47:17 2021 +0800

    ASoC: mediatek: mt8183: fix device_node leak

    [ Upstream commit cb006006fe6221f092fadaffd3f219288304c9ad ]

    Fixes the device_node leak.

    Signed-off-by: Tzung-Bi Shih <tzungbi@google.com>
    Link: https://lore.kernel.org/r/20211224064719.2031210-3-tzungbi@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b05f54e865d81fc2eb0d6ba010aa3b25e1a12dc6
Author: Tzung-Bi Shih <tzungbi@google.com>
Date:   Fri Dec 24 14:47:16 2021 +0800
```

```
    ASoC: mediatek: mt8173: fix device_node leak

    [ Upstream commit 493433785df0075afc0c106ab65f10a605d0b35d ]

    Fixes the device_node leak.

    Signed-off-by: Tzung-Bi Shih <tzungbi@google.com>
    Link: https://lore.kernel.org/r/20211224064719.2031210-2-tzungbi@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a85298b1ef4f35361288bbed769e73fcc9bb5d5
Author: Chunfeng Yun <chunfeng.yun@mediatek.com>
Date:   Sat Dec 18 16:27:59 2021 +0800

    phy: phy-mtk-tphy: add support efuse setting

    [ Upstream commit 6f2b033cb883f64ad084a75f13634242c7e179a6 ]

    Due to some SoCs have a bit shift issue that will drop a bit for usb3
    phy or pcie phy, fix it by adding software efuse reading and setting,
    but only support it optionally for version 2/3.

    Signed-off-by: Chunfeng Yun <chunfeng.yun@mediatek.com>
    Link: https://lore.kernel.org/r/20211218082802.5256-2-chunfeng.yun@mediatek.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1be8577f0b2f679095d237aaf281dca344f06c4
Author: Tzung-Bi Shih <tzungbi@google.com>
Date:   Tue Dec 14 12:00:28 2021 +0800

    ASoC: mediatek: mt8192-mt6359: fix device_node leak

    [ Upstream commit 4e28491a7a198c668437f2be8a91a76aa52f20eb ]

    The of_parse_phandle() document:
        >>> Use of_node_put() on it when done.

    The driver didn't call of_node_put().  Fixes the leak.

    Signed-off-by: Tzung-Bi Shih <tzungbi@google.com>
    Link: https://lore.kernel.org/r/20211214040028.2992627-1-tzungbi@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57f0460096c5854d99ce272341b083a66877dd05
Author: Sreekanth Reddy <sreekanth.reddy@broadcom.com>
Date:   Mon Dec 20 19:41:58 2021 +0530

    scsi: mpi3mr: Fixes around reply request queues

    [ Upstream commit 243bcc8efdb1f44b1a1d415e6821a246714c68ce ]

    Set reply queue depth of 1K for B0 and 4K for A0.

    While freeing the segmented request queues use the actual queue depth that
    is used while creating them.

    Link: https://lore.kernel.org/r/20211220141159.16117-25-sreekanth.reddy@broadcom.com
    Signed-off-by: Sreekanth Reddy <sreekanth.reddy@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8b4fd2c23ffefc47c43cbaffe0c0edd4bce7857a
Author: Christoph Hellwig <hch@lst.de>
Date:   Wed Dec 22 10:08:42 2021 +0100

    scsi: sr: Don't use GFP_DMA

    [ Upstream commit d94d94969a4ba07a43d62429c60372320519c391 ]

    The allocated buffers are used as a command payload, for which the block
    layer and/or DMA API do the proper bounce buffering if needed.

    Link: https://lore.kernel.org/r/20211222090842.920724-1-hch@lst.de
    Reported-by: Baoquan He <bhe@redhat.com>
    Reviewed-by: Baoquan He <bhe@redhat.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 43ce930314350567fcaa3b5aba3f064f408c3864
Author: Tianjia Zhang <tianjia.zhang@linux.alibaba.com>
Date:   Thu Dec 16 17:50:14 2021 +0800

    MIPS: Octeon: Fix build errors using clang

    [ Upstream commit 95339b70677dc6f9a2d669c4716058e71b8dc1c7 ]

    A large number of the following errors is reported when compiling
    with clang:

      cvmx-bootinfo.h:326:3: error: adding 'int' to a string does not append to the string
[-Werror,-Wstring-plus-int]
                      ENUM_BRD_TYPE_CASE(CVMX_BOARD_TYPE_NULL)
                      ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
      cvmx-bootinfo.h:321:20: note: expanded from macro 'ENUM_BRD_TYPE_CASE'
              case x: return(#x + 16);        /* Skip CVMX_BOARD_TYPE_ */
                             ~~~^~~~
      cvmx-bootinfo.h:326:3: note: use array indexing to silence this warning
      cvmx-bootinfo.h:321:20: note: expanded from macro 'ENUM_BRD_TYPE_CASE'
              case x: return(#x + 16);        /* Skip CVMX_BOARD_TYPE_ */
                              ^

    Follow the prompts to use the address operator '&' to fix this error.

    Signed-off-by: Tianjia Zhang <tianjia.zhang@linux.alibaba.com>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Philippe Mathieu-Daudé <f4bug@amsat.org>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 46f4ed654c150ae192efa7e4b3f3e218bdf2f099
Author: Michael Ellerman <mpe@ellerman.id.au>
Date:   Thu Dec 9 22:59:44 2021 +1100

    selftests/powerpc: Add a test of sigreturning to the kernel

    [ Upstream commit a8968521cfdc3e339fe69473d6632e0aa8d7202a ]

    We have a general signal fuzzer, sigfuz, which can modify the MSR & NIP
    before sigreturn. But the chance of it hitting a kernel address and also
    clearing MSR_PR is fairly slim.

    So add a specific test of sigreturn to a kernel address, both with and
    without attempting to clear MSR_PR (which the kernel must block).

    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
```

```
    Link: https://lore.kernel.org/r/20211209115944.4062384-1-mpe@ellerman.id.au
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd15ecbc16e2b52590da7fbac6ed819421e5b77e
Author: Ajit Kumar Pandey <AjitKumar.Pandey@amd.com>
Date:   Thu Dec 16 17:24:21 2021 -0600

    ASoC: SOF: ipc: Add null pointer check for substream->runtime

    [ Upstream commit 182b682b9ab1348e07ea1bf9d8f2505cc67f9237 ]

    When pcm stream is stopped "substream->runtime" pointer will be set
    to NULL by ALSA core. In case host received an ipc msg from firmware
    of type IPC_STREAM_POSITION after pcm stream is stopped, there will
    be kernel NULL pointer exception in ipc_period_elapsed(). This patch
    fixes it by adding NULL pointer check for "substream->runtime".

    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Ajit Kumar Pandey <AjitKumar.Pandey@amd.com>
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20211216232422.345164-3-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6efd329077ec67a829be915edfccb453cedb7a8
Author: Qi Liu <liuqi115@huawei.com>
Date:   Wed Dec 15 22:37:39 2021 +0800

    scsi: hisi_sas: Prevent parallel FLR and controller reset

    [ Upstream commit 16775db613c2bdea09705dcb876942c0641a1098 ]

    If we issue a controller reset command during executing a FLR a hung task
    may be found:

     Call trace:
      __switch_to+0x158/0x1cc
      __schedule+0x2e8/0x85c
      schedule+0x7c/0x110
      schedule_timeout+0x190/0x1cc
      __down+0x7c/0xd4
      down+0x5c/0x7c
      hisi_sas_task_exec+0x510/0x680 [hisi_sas_main]
      hisi_sas_queue_command+0x24/0x30 [hisi_sas_main]
      smp_execute_task_sg+0xf4/0x23c [libsas]
      sas_smp_phy_control+0x110/0x1e0 [libsas]
      transport_sas_phy_reset+0xc8/0x190 [libsas]
      phy_reset_work+0x2c/0x40 [libsas]
      process_one_work+0x1dc/0x48c
      worker_thread+0x15c/0x464
      kthread+0x160/0x170
      ret_from_fork+0x10/0x18

    This is a race condition which occurs when the FLR completes first.

    Here the host HISI_SAS_RESETTING_BIT flag out gets of sync as
    HISI_SAS_RESETTING_BIT is not always cleared with the hisi_hba.sem held, so
    now only set/unset HISI_SAS_RESETTING_BIT under hisi_hba.sem .

    Link: https://lore.kernel.org/r/1639579061-179473-7-git-send-email-john.garry@huawei.com
    Signed-off-by: Qi Liu <liuqi115@huawei.com>
    Signed-off-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 2992666f6005772d164a1505bf7982d322f09a8d
Author: Lakshmi Sowjanya D <lakshmi.sowjanya.d@intel.com>
Date:   Wed Dec 15 17:12:01 2021 +0200

    i2c: designware-pci: Fix to change data types of hcnt and lcnt parameters

    [ Upstream commit d52097010078c1844348dc0e467305e5f90fd317 ]

    The data type of hcnt and lcnt in the struct dw_i2c_dev is of type u16.
    It's better to have same data type in struct dw_scl_sda_cfg as well.

    Reported-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Lakshmi Sowjanya D <lakshmi.sowjanya.d@intel.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Jarkko Nikula <jarkko.nikula@linux.intel.com>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66f5aa73651981dbd43be2e4126be80135b4c66d
Author: Marc Zyngier <maz@kernel.org>
Date:   Thu Dec 16 14:32:27 2021 +0000

    irqchip/gic-v4: Disable redistributors' view of the VPE table at boot time

    [ Upstream commit 79a7f77b9b154d572bd9d2f1eecf58c4d018d8e2 ]

    Jay Chen reported that using a kdump kernel on a GICv4.1 system
    results in a RAS error being delivered when the secondary kernel
    configures the ITS's view of the new VPE table.

    As it turns out, that's because each RD still has a pointer to
    the previous instance of the VPE table, and that particular
    implementation is very upset by seeing two bits of the HW that
    should point to the same table with different values.

    To solve this, let's invalidate any reference that any RD has to
    the VPE table when discovering the RDs. The ITS can then be
    programmed as expected.

    Reported-by: Jay Chen <jkchen@linux.alibaba.com>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Cc: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Link: https://lore.kernel.org/r/20211214064716.21407-1-jkchen@linux.alibaba.com
    Link: https://lore.kernel.org/r/20211216144804.1578566-1-maz@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e7d4405bbb24cdbfc41143f6ed70d345d159009
Author: Ye Guojin <ye.guojin@zte.com.cn>
Date:   Tue Nov 16 08:10:51 2021 +0000

    MIPS: OCTEON: add put_device() after of_find_device_by_node()

    [ Upstream commit 858779df1c0787d3fec827fb705708df9ebdb15b ]

    This was found by coccicheck:
    ./arch/mips/cavium-octeon/octeon-platform.c, 332, 1-7, ERROR missing
    put_device; call of_find_device_by_node on line 324, but without a
    corresponding object release within this function.
    ./arch/mips/cavium-octeon/octeon-platform.c, 395, 1-7, ERROR missing
    put_device; call of_find_device_by_node on line 387, but without a
    corresponding object release within this function.
    ./arch/mips/cavium-octeon/octeon-usb.c, 512, 3-9, ERROR missing
    put_device; call of_find_device_by_node on line 515, but without a
    corresponding object release within this function.
```

```
      ./arch/mips/cavium-octeon/octeon-usb.c, 543, 1-7, ERROR missing
      put_device; call of_find_device_by_node on line 515, but without a
      corresponding object release within this function.

      Reported-by: Zeal Robot <zealci@zte.com.cn>
      Signed-off-by: Ye Guojin <ye.guojin@zte.com.cn>
      Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 899b8414230121d254c0e71002d242bd0f059e03
  Author: Jan Kara <jack@suse.cz>
  Date:   Tue Dec 14 11:04:29 2021 +0100

      udf: Fix error handling in udf_new_inode()

      [ Upstream commit f05f2429eec60851b98bdde213de31dab697c01b ]

      When memory allocation of iinfo or block allocation fails, already
      allocated struct udf_inode_info gets freed with iput() and
      udf_evict_inode() may look at inode fields which are not properly
      initialized. Fix it by marking inode bad before dropping reference to it
      in udf_new_inode().

      Reported-by: syzbot+9ca499bb57a2b9e4c652@syzkaller.appspotmail.com
      Signed-off-by: Jan Kara <jack@suse.cz>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit ede273eb4c007418b62e14f540a4408da6306a00
  Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
  Date:   Tue Dec 7 13:39:45 2021 -0600

      ASoC: SOF: Intel: hda: add quirks for HDAudio DMA position information

      [ Upstream commit 288fad2f71fa0b989c075d4984879c26d47cfb06 ]

      The code inherited from the Skylake driver does not seem to follow any
      known hardware recommendations.

      The only two recommended options are
      a) use DPIB registers if VC1 traffic is not allowed
      b) use DPIB DDR update if VC1 traffic is used

      In all of SOF-based updated, VC1 is not supported so we can 'safely'
      move to using DPIB registers only.

      This patch keeps the legacy code, in case there was an undocumented
      issue lost to history, and adds the DPIB DDR update for additional
      debug.

      Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
      Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
      Reviewed-by: Péter Ujfalusi <peter.ujfalusi@linux.intel.com>
      Link: https://lore.kernel.org/r/20211207193947.71080-6-pierre-
  louis.bossart@linux.intel.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 850a8dc610c287c2bac3073c376fde7155c7a7b6
  Author: Hari Bathini <hbathini@linux.ibm.com>
  Date:   Tue Dec 7 16:07:19 2021 +0530

      powerpc/fadump: Fix inaccurate CPU state info in vmcore generated with panic

      [ Upstream commit 06e629c25daa519be620a8c17359ae8fc7a2e903 ]
```

In panic path, fadump is triggered via a panic notifier function.
Before calling panic notifier functions, smp_send_stop() gets called,
which stops all CPUs except the panic'ing CPU. Commit 8389b37dffdc
("powerpc: stop_this_cpu: remove the cpu from the online map.") and
again commit bab26238bbd4 ("powerpc: Offline CPU in stop_this_cpu()")
started marking CPUs as offline while stopping them. So, if a kernel
has either of the above commits, vmcore captured with fadump via panic
path would not process register data for all CPUs except the panic'ing
CPU. Sample output of crash-utility with such vmcore:

```
  # crash vmlinux vmcore
  ...
        KERNEL: vmlinux
      DUMPFILE: vmcore  [PARTIAL DUMP]
          CPUS: 1
          DATE: Wed Nov 10 09:56:34 EST 2021
        UPTIME: 00:00:42
  LOAD AVERAGE: 2.27, 0.69, 0.24
         TASKS: 183
      NODENAME: XXXXXXXXX
       RELEASE: 5.15.0+
       VERSION: #974 SMP Wed Nov 10 04:18:19 CST 2021
       MACHINE: ppc64le  (2500 Mhz)
        MEMORY: 8 GB
         PANIC: "Kernel panic - not syncing: sysrq triggered crash"
           PID: 3394
       COMMAND: "bash"
          TASK: c0000000150a5f80  [THREAD_INFO: c0000000150a5f80]
           CPU: 1
         STATE: TASK_RUNNING (PANIC)

  crash> p -x __cpu_online_mask
  __cpu_online_mask = $1 = {
    bits = {0x2, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0}
  }
  crash>
  crash>
  crash> p -x __cpu_active_mask
  __cpu_active_mask = $2 = {
    bits = {0xff, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0, 0x0}
  }
  crash>
```

While this has been the case since fadump was introduced, the issue
was not identified for two probable reasons:

  - In general, the bulk of the vmcores analyzed were from crash
    due to exception.

  - The above did change since commit 8341f2f222d7 ("sysrq: Use
    panic() to force a crash") started using panic() instead of
    deferencing NULL pointer to force a kernel crash. But then
    commit de6e5d38417e ("powerpc: smp_send_stop do not offline
    stopped CPUs") stopped marking CPUs as offline till kernel
    commit bab26238bbd4 ("powerpc: Offline CPU in stop_this_cpu()")
    reverted that change.

To ensure post processing register data of all other CPUs happens
as intended, let panic() function take the crash friendly path (read
crash_smp_send_stop()) with the help of crash_kexec_post_notifiers
option. Also, as register data for all CPUs is captured by f/w, skip
IPI callbacks here for fadump, to avoid any complications in finding
the right backtraces.

Signed-off-by: Hari Bathini <hbathini@linux.ibm.com>

```
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211207103719.91117-2-hbathini@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f616a1bbdd0e1baca32480caa39093ab55e49329
Author: Hari Bathini <hbathini@linux.ibm.com>
Date:   Tue Dec 7 16:07:18 2021 +0530

    powerpc: handle kdump appropriately with crash_kexec_post_notifiers option

    [ Upstream commit 219572d2fc4135b5ce65c735d881787d48b10e71 ]

    Kdump can be triggered after panic_notifers since commit f06e5153f4ae2
    ("kernel/panic.c: add "crash_kexec_post_notifiers" option for kdump
    after panic_notifers") introduced crash_kexec_post_notifiers option.
    But using this option would mean smp_send_stop(), that marks all other
    CPUs as offline, gets called before kdump is triggered. As a result,
    kdump routines fail to save other CPUs' registers. To fix this, kdump
    friendly crash_smp_send_stop() function was introduced with kernel
    commit 0ee59413c967 ("x86/panic: replace smp_send_stop() with kdump
    friendly version in panic path"). Override this kdump friendly weak
    function to handle crash_kexec_post_notifiers option appropriately
    on powerpc.

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Hari Bathini <hbathini@linux.ibm.com>
    [Fixed signature of crash_stop_this_cpu() - reported by lkp@intel.com]
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211207103719.91117-1-hbathini@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6009ca52f886fa03d52eeccdb2c49399db999146
Author: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
Date:   Tue Dec 7 10:05:57 2021 -0300

    selftests/powerpc/spectre_v2: Return skip code when miss_percent is high

    [ Upstream commit 3c42e9542050d49610077e083c7c7c3f5fd5e26820 ]

    A mis-match between reported and actual mitigation is not restricted to the
    Vulnerable case. The guest might also report the mitigation as "Software
    count cache flush" and the host will still mitigate with branch cache
    disabled.

    So, instead of skipping depending on the detected mitigation, simply skip
    whenever the detected miss_percent is the expected one for a fully
    mitigated system, that is, above 95%.

    Signed-off-by: Thadeu Lima de Souza Cascardo <cascardo@canonical.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211207130557.40566-1-cascardo@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae00af6368846a69edf9d54a46c6e2df18e556c0
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Mon Sep 27 17:12:39 2021 +0200

    powerpc/40x: Map 32Mbytes of memory at startup

    [ Upstream commit 06e7cbc29e97b4713b4ea6def04ae8501a7d1a59 ]

    As reported by Carlo, 16Mbytes is not enough with modern kernels
    that tend to be a bit big, so map another 16M page at boot.

    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
```

```
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
 https://lore.kernel.org/r/89b5f974a7fa5011206682cd092e2c905530ff46.1632755552.git.christophe
 .leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 727054a761539009076ddf2192be986ca2b55bf8
 Author: Nathan Chancellor <nathan@kernel.org>
 Date:   Wed Dec 8 09:56:17 2021 -0700

    MIPS: Loongson64: Use three arguments for slti

    [ Upstream commit f2c6c22fa83ab2577619009057b3ebcb5305bb03 ]

    LLVM's integrated assembler does not support 'slti <reg>, <imm>':

    <instantiation>:16:12: error: invalid operand for instruction
     slti $12, (0x6300 | 0x0008)
              ^
    arch/mips/kernel/head.S:86:2: note: while in macro instantiation
     kernel_entry_setup # cpu specific setup
     ^
    <instantiation>:16:12: error: invalid operand for instruction
     slti $12, (0x6300 | 0x0008)
              ^
    arch/mips/kernel/head.S:150:2: note: while in macro instantiation
     smp_slave_setup
     ^

    To increase compatibility with LLVM's integrated assembler, use the full
    form of 'slti <reg>, <reg>, <imm>', which matches the rest of
    arch/mips/. This does not result in any change for GNU as.

    Link: https://github.com/ClangBuiltLinux/linux/issues/1526
    Reported-by: Ryutaroh Matsumoto <ryutaroh@ict.e.titech.ac.jp>
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 9ad2390d432543b34da5044604ae469e59007635
 Author: Takashi Iwai <tiwai@suse.de>
 Date:   Tue Dec 7 17:51:46 2021 +0100

    ALSA: seq: Set upper limit of processed events

    [ Upstream commit 6fadb494a638d8b8a55864ecc6ac58194f03f327 ]

    Currently ALSA sequencer core tries to process the queued events as
    much as possible when they become dispatchable.  If applications try
    to queue too massive events to be processed at the very same timing,
    the sequencer core would still try to process such all events, either
    in the interrupt context or via some notifier; in either away, it
    might be a cause of RCU stall or such problems.

    As a potential workaround for those problems, this patch adds the
    upper limit of the amount of events to be processed.  The remaining
    events are processed in the next batch, so they won't be lost.

    For the time being, it's limited up to 1000 events per queue, which
    should be high enough for any normal usages.

    Reported-by: Zqiang <qiang.zhang1211@gmail.com>
    Reported-by: syzbot+bb950e68b400ab4f65f8@syzkaller.appspotmail.com
    Link: https://lore.kernel.org/r/20211102033222.3849-1-qiang.zhang1211@gmail.com
    Link: https://lore.kernel.org/r/20211207165146.2888-1-tiwai@suse.de
```

```
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 841dad09ebeb3bb449987d0fa9d2e3e20f8c6f2f
Author: James Smart <jsmart2021@gmail.com>
Date:   Fri Dec 3 16:26:40 2021 -0800

    scsi: lpfc: Trigger SLI4 firmware dump before doing driver cleanup

    [ Upstream commit 7dd2e2a923173d637c272e483966be8e96a72b64 ]

    Extraneous teardown routines are present in the firmware dump path causing
    altered states in firmware captures.

    When a firmware dump is requested via sysfs, trigger the dump immediately
    without tearing down structures and changing adapter state.

    The driver shall rely on pre-existing firmware error state clean up
    handlers to restore the adapter.

    Link: https://lore.kernel.org/r/20211204002644.116455-6-jsmart2021@gmail.com
    Co-developed-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: James Smart <jsmart2021@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fc1ae9bb2a7ff574db3de1d01cd1daea1994ba14
Author: James Smart <jsmart2021@gmail.com>
Date:   Fri Dec 3 16:26:36 2021 -0800

    scsi: lpfc: Fix leaked lpfc_dmabuf mbox allocations with NPIV

    [ Upstream commit f0d3919697492950f57a26a1093aee53880d669d ]

    During rmmod testing, messages appeared indicating lpfc_mbuf_pool entries
    were still busy. This situation was only seen doing rmmod after at least 1
    vport (NPIV) instance was created and destroyed. The number of messages
    scaled with the number of vports created.

    When a vport is created, it can receive a PLOGI from another initiator
    Nport.  When this happens, the driver prepares to ack the PLOGI and
    prepares an RPI for registration (via mbx cmd) which includes an mbuf
    allocation. During the unsolicited PLOGI processing and after the RPI
    preparation, the driver recognizes it is one of the vport instances and
    decides to reject the PLOGI. During the LS_RJT preparation for the PLOGI,
    the mailbox struct allocated for RPI registration is freed, but the mbuf
    that was also allocated is not released.

    Fix by freeing the mbuf with the mailbox struct in the LS_RJT path.

    As part of the code review to figure the issue out a couple of other areas
    where found that also would not have released the mbuf. Those are cleaned
    up as well.

    Link: https://lore.kernel.org/r/20211204002644.116455-2-jsmart2021@gmail.com
    Co-developed-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: James Smart <jsmart2021@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 39f72d0af4a933b96925108367b52c252b923b1b
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Fri Dec 3 15:19:47 2021 -0800
```

```
    scsi: ufs: Fix a kernel crash during shutdown

    [ Upstream commit 3489c34bd02b73a72646037d673a122a53cee174 ]

    Fix the following kernel crash:

    Unable to handle kernel paging request at virtual address ffffffc91e735000
    Call trace:
     __queue_work+0x26c/0x624
     queue_work_on+0x6c/0xf0
     ufshcd_hold+0x12c/0x210
     __ufshcd_wl_suspend+0xc0/0x400
     ufshcd_wl_shutdown+0xb8/0xcc
     device_shutdown+0x184/0x224
     kernel_restart+0x4c/0x124
     __arm64_sys_reboot+0x194/0x264
     el0_svc_common+0xc8/0x1d4
     do_el0_svc+0x30/0x8c
     el0_svc+0x20/0x30
     el0_sync_handler+0x84/0xe4
     el0_sync+0x1bc/0x1c0

    Fix this crash by ungating the clock before destroying the work queue on
    which clock gating work is queued.

    Link: https://lore.kernel.org/r/20211203231950.193369-15-bvanassche@acm.org
    Tested-by: Bean Huo <beanhuo@micron.com>
    Reviewed-by: Bean Huo <beanhuo@micron.com>
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ef783bd64f63db3f156a2b4cdafb7022d97ad8d
Author: Stephan Gerhold <stephan@gerhold.net>
Date:   Mon Dec 6 12:45:42 2021 +0100

    interconnect: qcom: rpm: Prevent integer overflow in rate

    [ Upstream commit a7d9436a6c85fcb8843c910fd323dcd7f839bf63 ]

    Using icc-rpm on ARM32 currently results in clk_set_rate() errors during
    boot, e.g. "bus clk_set_rate error: -22". This is very similar to commit
    7381e27b1e56 ("interconnect: qcom: msm8974: Prevent integer overflow in rate")
    where the u64 is converted to a signed long during clock rate rounding,
    resulting in an overflow on 32-bit platforms.

    Let's fix it similarly by making sure that the rate does not exceed
    LONG_MAX. Such high clock rates will surely result in the maximum
    frequency of the bus anyway.

    Signed-off-by: Stephan Gerhold <stephan@gerhold.net>
    Link: https://lore.kernel.org/r/20211206114542.45325-1-stephan@gerhold.net
    Signed-off-by: Georgi Djakov <djakov@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6545406467dd185657d5c38e466430c7c207c279
Author: Ameer Hamza <amhamza.mgc@gmail.com>
Date:   Mon Dec 6 01:42:00 2021 +0500

    ASoC: test-component: fix null pointer dereference.

    [ Upstream commit c686316ec1210d43653c91e104c1e4cd0156dc89 ]

    Dereferncing of_id pointer will result in exception in current
```

```
      implementation since of_match_device() will assign it to NULL.
      Adding NULL check for protection.

      Signed-off-by: Ameer Hamza <amhamza.mgc@gmail.com>
      Link: https://lore.kernel.org/r/20211205204200.7852-1-amhamza.mgc@gmail.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 121390bee73f9cc3ab3ee039c1cd34b57a641e88
  Author: Christoph Hellwig <hch@lst.de>
  Date:   Mon Nov 29 11:21:36 2021 +0100

      dm: make the DAX support depend on CONFIG_FS_DAX

      [ Upstream commit 5d2a228b9e1319ff188f9ea89006fbe575561921 ]

      The device mapper DAX support is all hanging off a block device and thus
      can't be used with device dax.  Make it depend on CONFIG_FS_DAX instead
      of CONFIG_DAX_DRIVER.  This also means that bdev_dax_pgoff only needs to
      be built under CONFIG_FS_DAX now.

      Signed-off-by: Christoph Hellwig <hch@lst.de>
      Acked-by: Mike Snitzer <snitzer@redhat.com>
      Link: https://lore.kernel.org/r/20211129102203.2243509-3-hch@lst.de
      Signed-off-by: Dan Williams <dan.j.williams@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit a30798b751277e94b36414b40df4622f3a82e23e
  Author: Christoph Hellwig <hch@lst.de>
  Date:   Mon Nov 29 11:21:35 2021 +0100

      dm: fix alloc_dax error handling in alloc_dev

      [ Upstream commit d751939235b9b7bc4af15f90a3e99288a8b844a7 ]

      Make sure ->dax_dev is NULL on error so that the cleanup path doesn't
      trip over an ERR_PTR.

      Reported-by: Dan Williams <dan.j.williams@intel.com>
      Signed-off-by: Christoph Hellwig <hch@lst.de>
      Link: https://lore.kernel.org/r/20211129102203.2243509-2-hch@lst.de
      Signed-off-by: Dan Williams <dan.j.williams@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 2114bd1a316ffd31434fc543fd73b8778afc1d4e
  Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
  Date:   Tue Nov 30 13:39:09 2021 +0000

      nvmem: core: set size for sysfs bin file

      [ Upstream commit 86192251033308bb42f1e9813c962989d8ed07ec ]

      For some reason we never set the size for nvmem sysfs binary file.
      Set this.

      Reported-by: Gilles BULOZ <gilles.buloz@kontron.com>
      Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
      Link: https://lore.kernel.org/r/20211130133909.6154-1-srinivas.kandagatla@linaro.org
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 1eeabdd3677c29e40524a13b59b82ac4317630f0
  Author: Christophe Leroy <christophe.leroy@csgroup.eu>
  Date:   Fri Nov 26 18:06:46 2021 +0100
```

```
    w1: Misuse of get_user()/put_user() reported by sparse

    [ Upstream commit 33dc3e3e99e626ce51f462d883b05856c6c30b1d ]

    sparse warnings: (new ones prefixed by >>)
    >> drivers/w1/slaves/w1_ds28e04.c:342:13: sparse: sparse: incorrect type in initializer
(different address spaces) @@     expected char [noderef] __user *_pu_addr @@     got char
*buf @@
       drivers/w1/slaves/w1_ds28e04.c:342:13: sparse:     expected char [noderef] __user
*_pu_addr
       drivers/w1/slaves/w1_ds28e04.c:342:13: sparse:     got char *buf
    >> drivers/w1/slaves/w1_ds28e04.c:356:13: sparse: sparse: incorrect type in initializer
(different address spaces) @@     expected char const [noderef] __user *_gu_addr @@     got
char const *buf @@
       drivers/w1/slaves/w1_ds28e04.c:356:13: sparse:     expected char const [noderef]
__user *_gu_addr
       drivers/w1/slaves/w1_ds28e04.c:356:13: sparse:     got char const *buf

    The buffer buf is a failsafe buffer in kernel space, it's not user
    memory hence doesn't deserve the use of get_user() or put_user().

    Access 'buf' content directly.

    Link: https://lore.kernel.org/lkml/202111190526.K5vb7NWC-lkp@intel.com/T/
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Link:
https://lore.kernel.org/r/d14ed8d71ad4372e6839ae427f91441d3ba0e94d.1637946316.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bac5e46751d447125d51ccfeef58735fcfc40bc0
Author: Alexey Kardashevskiy <aik@ozlabs.ru>
Date:   Wed Sep 1 18:45:50 2021 +1000

    KVM: PPC: Book3S: Suppress failed alloc warning in H_COPY_TOFROM_GUEST

    [ Upstream commit 792020907b11c6f9246c21977cab3bad985ae4b6 ]

    H_COPY_TOFROM_GUEST is an hcall for an upper level VM to access its nested
    VMs memory. The userspace can trigger WARN_ON_ONCE(!(gfp & __GFP_NOWARN))
    in __alloc_pages() by constructing a tiny VM which only does
    H_COPY_TOFROM_GUEST with a too big GPR9 (number of bytes to copy).

    This silences the warning by adding __GFP_NOWARN.

    Spotted by syzkaller.

    Signed-off-by: Alexey Kardashevskiy <aik@ozlabs.ru>
    Reviewed-by: Fabiano Rosas <farosas@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20210901084550.1658699-1-aik@ozlabs.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 091a3826d9fc0d5cd236d633ec1c0ce460ea7839
Author: Alexey Kardashevskiy <aik@ozlabs.ru>
Date:   Wed Sep 1 18:45:12 2021 +1000

    KVM: PPC: Book3S: Suppress warnings when allocating too big memory slots

    [ Upstream commit 511d25d6b789fffcb20a3eb71899cf974a31bd9d ]

    The userspace can trigger "vmalloc size %lu allocation failure: exceeds
    total pages" via the KVM_SET_USER_MEMORY_REGION ioctl.
```

```
    This silences the warning by checking the limit before calling vzalloc()
    and returns ENOMEM if failed.

    This does not call underlying valloc helpers as __vmalloc_node() is only
    exported when CONFIG_TEST_VMALLOC_MODULE and __vmalloc_node_range() is
    not exported at all.

    Spotted by syzkaller.

    Signed-off-by: Alexey Kardashevskiy <aik@ozlabs.ru>
    [mpe: Use 'size' for the variable rather than 'cb']
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20210901084512.1658628-1-aik@ozlabs.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90748585b6459a2b03e5d759834ce635005c4aca
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Tue Nov 30 10:32:42 2021 +0100

    powerpc/powermac: Add missing lockdep_register_key()

    [ Upstream commit df1f679d19edb9eeb67cc2f96b29375f21991945 ]

    KeyWest i2c @0xf8001003 irq 42 /uni-n@f8000000/i2c@f8001000
    BUG: key c2d00cbc has not been registered!
    ------------[ cut here ]------------
    DEBUG_LOCKS_WARN_ON(1)
    WARNING: CPU: 0 PID: 1 at kernel/locking/lockdep.c:4801
lockdep_init_map_type+0x4c0/0xb4c
    Modules linked in:
    CPU: 0 PID: 1 Comm: swapper/0 Not tainted 5.15.5-gentoo-PowerMacG4 #9
    NIP:  c01a9428 LR: c01a9428 CTR: 00000000
    REGS: e1033cf0 TRAP: 0700   Not tainted  (5.15.5-gentoo-PowerMacG4)
    MSR:  00029032 <EE,ME,IR,DR,RI>  CR: 24002002  XER: 00000000

    GPR00: c01a9428 e1033db0 c2d1cf20 00000016 00000004 00000001 c01c0630 e1033a73
    GPR08: 00000000 00000000 00000000 e1033db0 24002004 00000000 f8729377 00000003
    GPR16: c1829a9c 00000000 18305357 c1416fc0 c1416f80 c006ac60 c2d00ca8 c1416f00
    GPR24: 00000000 c21586f0 c2160000 00000000 c2d00cbc c2170000 c216e1a0 c2160000
    NIP [c01a9428] lockdep_init_map_type+0x4c0/0xb4c
    LR [c01a9428] lockdep_init_map_type+0x4c0/0xb4c
    Call Trace:
    [e1033db0] [c01a9428] lockdep_init_map_type+0x4c0/0xb4c (unreliable)
    [e1033df0] [c1c177b8] kw_i2c_add+0x334/0x424
    [e1033e20] [c1c18294] pmac_i2c_init+0x9ec/0xa9c
    [e1033e80] [c1c1a790] smp_core99_probe+0xbc/0x35c
    [e1033eb0] [c1c03cb0] kernel_init_freeable+0x190/0x5a4
    [e1033f10] [c000946c] kernel_init+0x28/0x154
    [e1033f30] [c0035148] ret_from_kernel_thread+0x14/0x1c

    Add missing lockdep_register_key()

    Reported-by: Erhard Furtner <erhard_f@mailbox.org>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/69e4f55565bb45ebb0843977801b245af0c666fe.1638264741.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5831ecca63451a56be7cd7140576ba435256844f
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Sun Oct 31 14:50:06 2021 +0100
```

```
    clk: meson: gxbb: Fix the SDM_EN bit for MPLL0 on GXBB

    [ Upstream commit ff54938dd190d85f740b9bf9dde59b550936b621 ]

    There are reports that 48kHz audio does not work on the WeTek Play 2
    (which uses a GXBB SoC), while 44.1kHz audio works fine on the same
    board. There are also reports of 48kHz audio working fine on GXL and
    GXM SoCs, which are using an (almost) identical AIU (audio controller).

    Experimenting has shown that MPLL0 is causing this problem. In the .dts
    we have by default:
            assigned-clocks = <&clkc CLKID_MPLL0>,
                              <&clkc CLKID_MPLL1>,
                              <&clkc CLKID_MPLL2>;
            assigned-clock-rates = <294912000>,
                                   <270950400>,
                                   <393216000>;
    The MPLL0 rate is divisible by 48kHz without remainder and the MPLL1
    rate is divisible by 44.1kHz without remainder. Swapping these two clock
    rates "fixes" 48kHz audio but breaks 44.1kHz audio.

    Everything looks normal when looking at the info provided by the common
    clock framework while playing 48kHz audio (via I2S with mclk-fs = 256):
            mpll_prediv                 1        1        0  2000000000
               mpll0_div                1        1        0   294909641
                  mpll0                 1        1        0   294909641
                     cts_amclk_sel       1        1        0   294909641
                        cts_amclk_div       1        1        0    12287902
                           cts_amclk       1        1        0    12287902

    meson-clk-msr however shows that the actual MPLL0 clock is off by more
    than 38MHz:
            mp0_out               333322917    +/-10416Hz

    The rate seen by meson-clk-msr is very close to what we would get when
    SDM (the fractional part) was ignored:
      (2000000000Hz * 16384) / ((16384 * 6) = 333.33MHz
    If SDM was considered the we should get close to:
      (2000000000Hz * 16384) / ((16384 * 6) + 12808) = 294.9MHz

    Further experimenting shows that HHI_MPLL_CNTL7[15] does not have any
    effect on the rate of MPLL0 as seen my meson-clk-msr (regardless of
    whether that bit is zero or one the rate is always the same according to
    meson-clk-msr). Using HHI_MPLL_CNTL[25] on the other hand as SDM_EN
    results in SDM being considered for the rate output by the hardware. The
    rate - as seen by meson-clk-msr - matches with what we expect when
    SDM_EN is enabled (fractional part is being considered, resulting in a
    294.9MHz output) or disable (fractional part being ignored, resulting in
    a 333.33MHz output).

    Reported-by: Christian Hewitt <christianshewitt@gmail.com>
    Tested-by: Christian Hewitt <christianshewitt@gmail.com>
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Jerome Brunet <jbrunet@baylibre.com>
    Link: https://lore.kernel.org/r/20211031135006.1508796-1-
martin.blumenstingl@googlemail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a80843b69f14310af8830104bfe490e3542882b6
Author: Joakim Tjernlund <joakim.tjernlund@infinera.com>
Date:   Thu May 11 14:20:33 2017 +0200

    i2c: mpc: Correct I2C reset procedure

    [ Upstream commit ebe82cf92cd4825c3029434cabfcd2f1780e64be ]
```

```
    Current I2C reset procedure is broken in two ways:
    1) It only generate 1 START instead of 9 STARTs and STOP.
    2) It leaves the bus Busy so every I2C xfer after the first
       fixup calls the reset routine again, for every xfer there after.

    This fixes both errors.

    Signed-off-by: Joakim Tjernlund <joakim.tjernlund@infinera.com>
    Acked-by: Scott Wood <oss@buserror.net>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35ba6505ff4ab8c3b7c65f2bcf198c7c3cbf0e37
Author: Michael Ellerman <mpe@ellerman.id.au>
Date:   Wed Nov 24 20:32:53 2021 +1100

    powerpc/smp: Move setup_profiling_timer() under CONFIG_PROFILING

    [ Upstream commit a4ac0d249a5db80e79d573db9e4ad29354b643a8 ]

    setup_profiling_timer() is only needed when CONFIG_PROFILING is enabled.

    Fixes the following W=1 warning when CONFIG_PROFILING=n:
      linux/arch/powerpc/kernel/smp.c:1638:5: error: no previous prototype for
‘setup_profiling_timer’

    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211124093254.1054750-5-mpe@ellerman.id.au
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07d9beb6e3c2e852e884113d6803ea4b3643ae38
Author: Heiner Kallweit <hkallweit1@gmail.com>
Date:   Sun Nov 7 22:57:00 2021 +0100

    i2c: i801: Don't silently correct invalid transfer size

    [ Upstream commit effa453168a7eeb8a562ff4edc1dbf9067360a61 ]

    If an invalid block size is provided, reject it instead of silently
    changing it to a supported value. Especially critical I see the case of
    a write transfer with block length 0. In this case we have no guarantee
    that the byte we would write is valid. When silently reducing a read to
    32 bytes then we don't return an error and the caller may falsely
    assume that we returned the full requested data.

    If this change should break any (broken) caller, then I think we should
    fix the caller.

    Signed-off-by: Heiner Kallweit <hkallweit1@gmail.com>
    Reviewed-by: Jean Delvare <jdelvare@suse.de>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63d3db4c0d5e0ee3ae4459dec3209b7af93c6454
Author: Ye Guojin <ye.guojin@zte.com.cn>
Date:   Wed Nov 10 00:29:10 2021 +0000

    ASoC: imx-hdmi: add put_device() after of_find_device_by_node()

    [ Upstream commit f670b274f7f6f4b2722d7f08d0fddf606a727e92 ]

    This was found by coccicheck:
    ./sound/soc/fsl/imx-hdmi.c,209,1-7,ERROR  missing put_device; call
    of_find_device_by_node on line 119, but without a corresponding object
```

```
    release within this function.

    Reported-by: Zeal Robot <zealci@zte.com.cn>
    Signed-off-by: Ye Guojin <ye.guojin@zte.com.cn>
    Link: https://lore.kernel.org/r/20211110002910.134915-1-ye.guojin@zte.com.cn
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1f637bb828edd71e0a6ec793d445eaa75e98856
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Wed Nov 10 12:50:53 2021 +1000

    powerpc/watchdog: Fix missed watchdog reset due to memory ordering race

    [ Upstream commit 5dad4ba68a2483fc80d70b9dc90bbe16e1f27263 ]

    It is possible for all CPUs to miss the pending cpumask becoming clear,
    and then nobody resetting it, which will cause the lockup detector to
    stop working. It will eventually expire, but watchdog_smp_panic will
    avoid doing anything if the pending mask is clear and it will never be
    reset.

    Order the cpumask clear vs the subsequent test to close this race.

    Add an extra check for an empty pending mask when the watchdog fires and
    finds its bit still clear, to try to catch any other possible races or
    bugs here and keep the watchdog working. The extra test in
    arch_touch_nmi_watchdog is required to prevent the new warning from
    firing off.

    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Reviewed-by: Laurent Dufour <ldufour@linux.ibm.com>
    Debugged-by: Laurent Dufour <ldufour@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211110025056.2084347-2-npiggin@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01de442c0721acc40dab34ed1dead3de4287a6cf
Author: Julia Lawall <Julia.Lawall@lip6.fr>
Date:   Fri Nov 20 20:33:23 2015 +0000

    powerpc/btext: add missing of_node_put

    [ Upstream commit a1d2b210ffa52d60acabbf7b6af3ef7e1e69cda0 ]

    for_each_node_by_type performs an of_node_get on each iteration, so
    a break out of the loop requires an of_node_put.

    A simplified version of the semantic patch that fixes this problem is as
    follows (http://coccinelle.lip6.fr):

    // <smpl>
    @@
    local idexpression n;
    expression e;
    @@

     for_each_node_by_type(n,...) {
       ...
    (
       of_node_put(n);
    |
       e = n
    |
    +  of_node_put(n);
```

```
     ?  break;
     )
        ...
      }
     ... when != n
     // </smpl>

     Signed-off-by: Julia Lawall <Julia.Lawall@lip6.fr>
     Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
     Link: https://lore.kernel.org/r/1448051604-25256-6-git-send-email-Julia.Lawall@lip6.fr
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fbedd0b762fa73fba2f6019d1812f7bb4ea2afce
  Author: Julia Lawall <Julia.Lawall@lip6.fr>
  Date:   Fri Nov 20 21:33:24 2015 +0100

     powerpc/cell: add missing of_node_put

     [ Upstream commit a841fd009e51c8c0a8f07c942e9ab6bb48da8858 ]

     for_each_node_by_name performs an of_node_get on each iteration, so
     a break out of the loop requires an of_node_put.

     A simplified version of the semantic patch that fixes this problem is as
     follows (http://coccinelle.lip6.fr):

     // <smpl>
     @@
     expression e,e1;
     local idexpression n;
     @@

      for_each_node_by_name(n, e1) {
        ... when != of_node_put(n)
            when != e = n
     (
        return n;
     |
     +  of_node_put(n);
     ?  return ...;
     )
        ...
      }
     // </smpl>

     Signed-off-by: Julia Lawall <Julia.Lawall@lip6.fr>
     Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
     Link: https://lore.kernel.org/r/1448051604-25256-7-git-send-email-Julia.Lawall@lip6.fr
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 5e1631287bd21335eeffefd9407161aaf6061944
  Author: Julia Lawall <Julia.Lawall@lip6.fr>
  Date:   Fri Nov 20 20:33:21 2015 +0000

     powerpc/powernv: add missing of_node_put

     [ Upstream commit 7d405a939ca960162eb30c1475759cb2fdf38f8c ]

     for_each_compatible_node performs an of_node_get on each iteration, so
     a break out of the loop requires an of_node_put.

     A simplified version of the semantic patch that fixes this problem is as
     follows (http://coccinelle.lip6.fr):

     // <smpl>
```

```
      @@
      local idexpression n;
      expression e;
      @@

       for_each_compatible_node(n,...) {
          ...
      (
          of_node_put(n);
      |
          e = n
      |
      +   of_node_put(n);
      ?   break;
      )
          ...
       }
      ... when != n
      // </smpl>

      Signed-off-by: Julia Lawall <Julia.Lawall@lip6.fr>
      Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
      Link: https://lore.kernel.org/r/1448051604-25256-4-git-send-email-Julia.Lawall@lip6.fr
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fb8191749f461973952142637249aa0c200f76a9
  Author: Julia Lawall <Julia.Lawall@lip6.fr>
  Date:   Fri Nov 20 20:33:19 2015 +0000

      powerpc/6xx: add missing of_node_put

      [ Upstream commit f6e82647ff71d427d4148964b71f239fba9d7937 ]

      for_each_compatible_node performs an of_node_get on each iteration, so
      a break out of the loop requires an of_node_put.

      A simplified version of the semantic patch that fixes this problem is as
      follows (http://coccinelle.lip6.fr):

      // <smpl>
      @@
      expression e;
      local idexpression n;
      @@

      @@
      local idexpression n;
      expression e;
      @@

       for_each_compatible_node(n,...) {
          ...
      (
          of_node_put(n);
      |
          e = n
      |
      +   of_node_put(n);
      ?   break;
      )
          ...
       }
      ... when != n
      // </smpl>
```

```
    Signed-off-by: Julia Lawall <Julia.Lawall@lip6.fr>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/1448051604-25256-2-git-send-email-Julia.Lawall@lip6.fr
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6755949142042a65792577d32157156e5cd9b375
Author: Sam Protsenko <semen.protsenko@linaro.org>
Date:   Mon Nov 22 16:42:06 2021 +0200

    clk: samsung: exynos850: Register clocks early

    [ Upstream commit bcda841f9bf2cddcf2f000cba96f2e27f6f2bdbf ]

    Some clocks must be registered before init calls. For example MCT clock
    (from CMU_PERI) is needed for MCT timer driver, which is registered
    with TIMER_OF_DECLARE(). By the time we get to core_initcall() used for
    clk-exynos850 platform driver init, it's already too late. Inability to
    get "mct" clock in MCT driver leads to kernel panic, as functions
    registered with *_OF_DECLARE() can't do deferred calls. MCT timer driver
    can't be fixed either, as it's acting as a clock source and it's
    essential to register it in start_kernel() -> time_init().

    Let's register CMU_PERI clocks early, using CLK_OF_DECLARE(). CMU_TOP
    generates clocks needed for CMU_PERI, but it's already registered early.

    While at it, let's cleanup the code a bit, by extracting everything
    related to CMU initialization and registration to the separate function.

    Similar issue was discussed at [1] and addressed in commit 1f7db7bbf031
    ("clk: renesas: cpg-mssr: Add early clock support"), as well as in
    drivers/clk/mediatek/clk-mt2712.c.

    [1] https://patchwork.kernel.org/project/linux-renesas-soc/patch/20180829132954.64862-2-
chris.brandt@renesas.com/

    Signed-off-by: Sam Protsenko <semen.protsenko@linaro.org>
    Signed-off-by: Sylwester Nawrocki <s.nawrocki@samsung.com>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20211122144206.23134-1-semen.protsenko@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fffd5d844ac2238b733ba352818d426c77525176
Author: Ingo Molnar <mingo@kernel.org>
Date:   Wed Jan 5 01:35:58 2022 +0100

    x86/kbuild: Enable CONFIG_KALLSYMS_ALL=y in the defconfigs

    [ Upstream commit b6aa86cff44cf099299d3a5e66348cb709cd7964 ]

    Most distro kernels have this option enabled, to improve debug output.

    Lockdep also selects it.

    Enable this in the defconfig kernel as well, to make it more
    representative of what people are using on x86.

    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Link: https://lore.kernel.org/r/YdTn7gssoMVDMgMw@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2614cccd17e6e7900693734f49e54f5db1acacd2
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Fri Jan 7 14:29:15 2022 +0100

    can: flexcan: add more quirks to describe RX path capabilities
```

```
    [ Upstream commit c5c88591040ee7d84d037328eed9019d3ffab821 ]

    Most flexcan IP cores support 2 RX modes:
    - FIFO
    - mailbox

    Some IP core versions cannot receive CAN RTR messages via mailboxes.
    This patch adds quirks to document this.

    This information will be used in a later patch to switch from FIFO to
    more performant mailbox mode at the expense of losing the ability to
    receive RTR messages. This trade off is beneficial in certain use
    cases.

    Link: https://lore.kernel.org/all/20220107193105.1699523-5-mkl@pengutronix.de
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d7358f4ba2af63407f9a18d96cc70de49dc3a410
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Thu Jan 6 12:45:46 2022 +0100

    can: flexcan: rename RX modes

    [ Upstream commit 34ea4e1c99f1f177f87e4ae7896caef238dd741a ]

    Most flexcan IP cores support 2 RX modes:
    - FIFO
    - mailbox

    The names for these modes were chosen to reflect the name of the
    rx-offload mode they are using.

    The name of the RX modes should better reflect their difference with
    regards the flexcan IP core. So this patch renames the various
    occurrences of OFF_FIFO to RX_FIFO and OFF_TIMESTAMP to RX_MAILBOX:

    | FLEXCAN_TX_MB_RESERVED_OFF_FIFO -> FLEXCAN_TX_MB_RESERVED_RX_FIFO
    | FLEXCAN_TX_MB_RESERVED_OFF_TIMESTAMP -> FLEXCAN_TX_MB_RESERVED_RX_MAILBOX
    | FLEXCAN_QUIRK_USE_OFF_TIMESTAMP -> FLEXCAN_QUIRK_USE_RX_MAILBOX

    Link: https://lore.kernel.org/all/20220107193105.1699523-4-mkl@pengutronix.de
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7eedb6caa0648f9ffcd26b4c16371841b6e8c40e
Author: Dario Binacchi <dario.binacchi@amarulasolutions.com>
Date:   Tue Jan 4 14:20:25 2022 +0100

    can: flexcan: allow to change quirks at runtime

    [ Upstream commit 01bb4dccd92b4dc21f6af3312e5696924e371111 ]

    This is a preparation patch for the upcoming support to change the
    rx-rtr capability via the ethtool API.

    Link: https://lore.kernel.org/all/20220107193105.1699523-3-mkl@pengutronix.de
    Signed-off-by: Dario Binacchi <dario.binacchi@amarulasolutions.com>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 228838dc4d5513750fb144867a013b77ac82f626
Author: John David Anglin <dave.anglin@bell.net>
Date:   Wed Dec 22 16:52:26 2021 +0000
```

```
    parisc: Avoid calling faulthandler_disabled() twice

    [ Upstream commit 9e9d4b460f23bab61672eae397417d03917d116c ]

    In handle_interruption(), we call faulthandler_disabled() to check whether the
    fault handler is not disabled. If the fault handler is disabled, we immediately
    call do_page_fault(). It then calls faulthandler_disabled(). If disabled,
    do_page_fault() attempts to fixup the exception by jumping to no_context:

    no_context:

            if (!user_mode(regs) && fixup_exception(regs)) {
                    return;
            }

            parisc_terminate("Bad Address (null pointer deref?)", regs, code, address);

    Apart from the error messages, the two blocks of code perform the same
    function.

    We can avoid two calls to faulthandler_disabled() by a simple revision
    to the code in handle_interruption().

    Note: I didn't try to fix the formatting of this code block.

    Signed-off-by: John David Anglin <dave.anglin@bell.net>
    Signed-off-by: Helge Deller <deller@gmx.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97b890eae18c9fdca2ab5ec1e51a21b268032b04
Author: Maor Dickman <maord@nvidia.com>
Date:   Thu Dec 9 14:03:01 2021 +0200

    net/mlx5e: Unblock setting vid 0 for VF in case PF isn't eswitch manager

    [ Upstream commit 7846665d3504812acaebf920d1141851379a7f37 ]

    When using libvirt to passthrough VF to VM it will always set the VF vlan
    to 0 even if user didn’t request it, this will cause libvirt to fail to
    boot in case the PF isn't eswitch owner.

    Example of such case is the DPU host PF which isn't eswitch manager, so
    any attempt to passthrough VF of it using libvirt will fail.

    Fix it by not returning error in case set VF vlan is called with vid 0.

    Signed-off-by: Maor Dickman <maord@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47c1676f9059c696fcd92db082a5a27b47cb4867
Author: Maher Sanalla <msanalla@nvidia.com>
Date:   Wed Jan 5 14:50:11 2022 +0200

    net/mlx5: Update log_max_qp value to FW max capability

    [ Upstream commit f79a609ea6bf54ad2d2c24e4de4524288b221666 ]

    log_max_qp in driver's default profile #2 was set to 18, but FW actually
    supports 17 at the most - a situation that led to the concerning print
    when the driver is loaded:
    "log_max_qp value in current profile is 18, changing to HCA capabaility
    limit (17)"
```

```
    The expected behavior from mlx5_profile #2 is to match the maximum FW
    capability in regards to log_max_qp. Thus, log_max_qp in profile #2 is
    initialized to a defined static value (0xff) - which basically means that
    when loading this profile, log_max_qp value  will be what the currently
    installed FW supports at most.

    Signed-off-by: Maher Sanalla <msanalla@nvidia.com>
    Reviewed-by: Maor Gottlieb <maorg@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f60df29fe0379b80ddcba7bc6991b36451d9d69
Author: Jason A. Donenfeld <Jason@zx2c4.com>
Date:   Wed Dec 29 22:10:05 2021 +0100

    random: do not throw away excess input to crng_fast_load

    [ Upstream commit 73c7733f122e8d0107f88655a12011f68f69e74b ]

    When crng_fast_load() is called by add_hwgenerator_randomness(), we
    currently will advance to crng_init==1 once we've acquired 64 bytes, and
    then throw away the rest of the buffer. Usually, that is not a problem:
    When add_hwgenerator_randomness() gets called via EFI or DT during
    setup_arch(), there won't be any IRQ randomness. Therefore, the 64 bytes
    passed by EFI exactly matches what is needed to advance to crng_init==1.
    Usually, DT seems to pass 64 bytes as well -- with one notable exception
    being kexec, which hands over 128 bytes of entropy to the kexec'd kernel.
    In that case, we'll advance to crng_init==1 once 64 of those bytes are
    consumed by crng_fast_load(), but won't continue onward feeding in bytes
    to progress to crng_init==2. This commit fixes the issue by feeding
    any leftover bytes into the next phase in add_hwgenerator_randomness().

    [linux@dominikbrodowski.net: rewrite commit message]
    Signed-off-by: Dominik Brodowski <linux@dominikbrodowski.net>
    Signed-off-by: Jason A. Donenfeld <Jason@zx2c4.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88d46d54dd921cdce2a0254f6e376bfae37cd081
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 2 18:52:44 2022 +0100

    serial: core: Keep mctrl register state and cached copy in sync

    [ Upstream commit 93a770b7e16772530196674ffc79bb13fa927dc6 ]

    struct uart_port contains a cached copy of the Modem Control signals.
    It is used to skip register writes in uart_update_mctrl() if the new
    signal state equals the old signal state.  It also avoids a register
    read to obtain the current state of output signals.

    When a uart_port is registered, uart_configure_port() changes signal
    state but neglects to keep the cached copy in sync.  That may cause
    a subsequent register write to be incorrectly skipped.  Fix it before
    it trips somebody up.

    This behavior has been present ever since the serial core was introduced
    in 2002:
    https://git.kernel.org/history/history/c/33c0d1b0c3eb

    So far it was never an issue because the cached copy is initialized to 0
    by kzalloc() and when uart_configure_port() is executed, at most DTR has
    been set by uart_set_options() or sunsu_console_setup().  Therefore,
    a stable designation seems unnecessary.
```

```
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Link:
https://lore.kernel.org/r/bceeaba030b028ed810272d55d5fc6f3656ddddb.1641129752.git.lukas@wunn
er.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 739427ae25881239d495a360c284752d6f4718f6
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 2 18:45:44 2022 +0100

    serial: pl011: Drop CR register reset on set_termios

    [ Upstream commit e368cc656fd6d0075f1c3ab9676e2001451e3e04 ]

    pl011_set_termios() briefly resets the CR register to zero, thereby
    glitching DTR/RTS signals.  With rs485 this may result in the bus being
    occupied for no reason.

    Where does this register write originate from?

    The PL011 driver was forked from the PL010 driver in 2004:
    https://git.kernel.org/history/history/c/157c0342e591

    Until this commit, the PL010 driver's IRQ handler ambauart_int()
    modified the CR register without holding the port spinlock.

    ambauart_set_termios() also modified that register.  To prevent
    concurrent read-modify-writes by the IRQ handler and to prevent
    transmission while changing baudrate, ambauart_set_termios() had to
    disable interrupts.  On the PL010, that is achieved by writing zero to
    the CR register.

    However, on the PL011, interrupts are disabled in the IMSC register,
    not in the CR register.

    Additionally, the commit amended both the PL010 and PL011 driver to
    acquire the port spinlock in the IRQ handler, obviating the need to
    disable interrupts in ->set_termios().

    So the CR register write is obsolete for two reasons.  Drop it.

    Cc: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Link:
https://lore.kernel.org/r/f49f945375f5ccb979893c49f1129f51651ac738.1641129062.git.lukas@wunn
er.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1fab42f4481ff7618876c6b00ea9d64bda23984
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sun Jan 2 18:42:44 2022 +0100

    serial: pl010: Drop CR register reset on set_termios

    [ Upstream commit 08a0c6dff91c965e39905cf200d22db989203ccb ]

    pl010_set_termios() briefly resets the CR register to zero.

    Where does this register write come from?

    The PL010 driver's IRQ handler ambauart_int() originally modified the CR
    register without holding the port spinlock.  ambauart_set_termios() also
    modified that register.  To prevent concurrent read-modify-writes by the
```

```
    IRQ handler and to prevent transmission while changing baudrate,
    ambauart_set_termios() had to disable interrupts.  That is achieved by
    writing zero to the CR register.

    However in 2004 the PL010 driver was amended to acquire the port
    spinlock in the IRQ handler, obviating the need to disable interrupts in
    ->set_termios():
    https://git.kernel.org/history/history/c/157c0342e591

    That rendered the CR register write obsolete.  Drop it.

    Cc: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Link:
https://lore.kernel.org/r/fcaff16e5b1abb4cc3da5a2879ac13f278b99ed0.1641128728.git.lukas@wunn
er.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9e107940cc4d551b054ac0f5cf0b0b83130eda14
Author: Konrad Dybcio <konrad.dybcio@somainline.org>
Date:   Thu Dec 30 03:34:42 2021 +0100

    regulator: qcom_smd: Align probe function with rpmh-regulator

    [ Upstream commit 14e2976fbabdacb01335d7f91eeebbc89c67ddb1 ]

    The RPMh regulator driver is much newer and gets more attention, which in
    consequence makes it do a few things better. Update qcom_smd-regulator's
    probe function to mimic what rpmh-regulator does to address a couple of
    issues:

    - Probe defer now works correctly, before it used to, well,
      kinda just die.. This fixes reliable probing on (at least) PM8994,
      because Linux apparently cannot deal with supply map dependencies yet..

    - Regulator data is now matched more sanely: regulator data is matched
      against each individual regulator node name and throwing an -EINVAL if
      data is missing, instead of just assuming everything is fine and
      iterating over all subsequent array members.

    - status = "disabled" will now work for disabling individual regulators in
      DT. Previously it didn't seem to do much if anything at all.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@somainline.org>
    Link: https://lore.kernel.org/r/20211230023442.1123424-1-konrad.dybcio@somainline.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eaac4434322748dd2e9ace2764e95c2aa4e17140
Author: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
Date:   Tue Jan 4 16:38:31 2022 +0000

    net: gemini: allow any RGMII interface mode

    [ Upstream commit 4e4f325a0a55907b14f579e6b1a38c53755e3de2 ]

    The four RGMII interface modes take care of the required RGMII delay
    configuration at the PHY and should not be limited by the network MAC
    driver. Sadly, gemini was only permitting RGMII mode with no delays,
    which would require the required delay to be inserted via PCB tracking
    or by the MAC.

    However, there are designs that require the PHY to add the delay, which
    is impossible without Gemini permitting the other three PHY interface
```

```
    modes. Fix the driver to allow these.

    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Tested-by: Corentin Labbe <clabbe.montjoie@gmail.com>
    Link: https://lore.kernel.org/r/E1n4mpT-002PLd-Ha@rmk-PC.armlinux.org.uk
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 596b4d50867c1cd7b684cba1b20df18fed9373f6
Author: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
Date:   Tue Jan 4 16:38:19 2022 +0000

    net: phy: marvell: configure RGMII delays for 88E1118

    [ Upstream commit f22725c95ececb703c3f741e8f946d23705630b7 ]

    Corentin Labbe reports that the SSI 1328 does not work when allowing
    the PHY to operate at gigabit speeds, but does work with the generic
    PHY driver.

    This appears to be because m88e1118_config_init() writes a fixed value
    to the MSCR register, claiming that this is to enable 1G speeds.
    However, this always sets bits 4 and 5, enabling RGMII transmit and
    receive delays. The suspicion is that the original board this was
    added for required the delays to make 1G speeds work.

    Add the necessary configuration for RGMII delays for the 88E1118 to
    bring this into line with the requirements for RGMII support, and thus
    make the SSI 1328 work.

    Corentin Labbe has tested this on gemini-ssi1328 and gemini-ns2502.

    Reported-by: Corentin Labbe <clabbe.montjoie@gmail.com>
    Tested-by: Corentin Labbe <clabbe.montjoie@gmail.com>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d386a80e75a01e7761f2edaa0088f5082ce393df
Author: Danielle Ratson <danieller@nvidia.com>
Date:   Wed Jan 5 12:22:27 2022 +0200

    mlxsw: pci: Avoid flow control for EMAD packets

    [ Upstream commit d43e4271747ace01a27a49a97a397cb4219f6487 ]

    Locally generated packets ingress the device through its CPU port. When
    the CPU port is congested and there are not enough credits in its
    headroom buffer, packets can be dropped.

    While this might be acceptable for data packets that traverse the
    network, configuration packets exchanged between the host and the device
    (EMADs) should not be subjected to this flow control.

    The "sdq_lp" bit in the SDQ (Send Descriptor Queue) context allows the
    host to instruct the device to treat packets sent on this queue as
    "local processing" and always process them, regardless of the state of
    the CPU port's headroom.

    Add the definition of this bit and set it for the dedicated SDQ reserved
    for the transmission of EMAD packets. This makes the "local processing"
    bit in the WQE (Work Queue Element) redundant, so clear it.
```

```
    Signed-off-by: Danielle Ratson <danieller@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 71b3ae7e103c3273cdeb5c80b6222d145d061bf8
Author: Jiri Olsa <jolsa@redhat.com>
Date:   Tue Jan 4 13:10:30 2022 +0100

    bpf/selftests: Fix namespace mount setup in tc_redirect

    [ Upstream commit 5e22dd18626726028a93ff1350a8a71a00fd843d ]

    The tc_redirect umounts /sys in the new namespace, which can be
    mounted as shared and cause global umount. The lazy umount also
    takes down mounted trees under /sys like debugfs, which won't be
    available after sysfs mounts again and could cause fails in other
    tests.

      # cat /proc/self/mountinfo | grep debugfs
      34 23 0:7 / /sys/kernel/debug rw,nosuid,nodev,noexec,relatime shared:14 - debugfs
debugfs rw
      # cat /proc/self/mountinfo | grep sysfs
      23 86 0:22 / /sys rw,nosuid,nodev,noexec,relatime shared:2 - sysfs sysfs rw
      # mount | grep debugfs
      debugfs on /sys/kernel/debug type debugfs (rw,nosuid,nodev,noexec,relatime)

      # ./test_progs -t tc_redirect
      #164 tc_redirect:OK
      Summary: 1/4 PASSED, 0 SKIPPED, 0 FAILED

      # mount | grep debugfs
      # cat /proc/self/mountinfo | grep debugfs
      # cat /proc/self/mountinfo | grep sysfs
      25 86 0:22 / /sys rw,relatime shared:2 - sysfs sysfs rw

    Making the sysfs private under the new namespace so the umount won't
    trigger the global sysfs umount.

    Reported-by: Hangbin Liu <haliu@redhat.com>
    Signed-off-by: Jiri Olsa <jolsa@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Cc: Jussi Maki <joamaki@gmail.com>
    Link: https://lore.kernel.org/bpf/20220104121030.138216-1-jolsa@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c46ea9333993873a1dc074b615c6d6bc8bd2ba6
Author: Vincent Mailhol <mailhol.vincent@wanadoo.fr>
Date:   Tue Dec 7 21:15:27 2021 +0900

    can: do not increase rx statistics when generating a CAN rx error message frame

    [ Upstream commit 676068db69b847f06fe054fca15bf6b107bd24da ]

    The CAN error message frames (i.e. error skb) are an interface
    specific to socket CAN. The payload of the CAN error message frames
    does not correspond to any actual data sent on the wire. Only an error
    flag and a delimiter are transmitted when an error occurs (c.f. ISO
    11898-1 section 10.4.4.2 "Error flag").

    For this reason, it makes no sense to increment the rx_packets and
    rx_bytes fields of struct net_device_stats because no actual payload
    were transmitted on the wire.

    This patch fixes all the CAN drivers.
```

```
    Link: https://lore.kernel.org/all/20211207121531.42941-2-mailhol.vincent@wanadoo.fr
    CC: Marc Kleine-Budde <mkl@pengutronix.de>
    CC: Nicolas Ferre <nicolas.ferre@microchip.com>
    CC: Alexandre Belloni <alexandre.belloni@bootlin.com>
    CC: Ludovic Desroches <ludovic.desroches@microchip.com>
    CC: Chandrasekar Ramakrishnan <rcsekar@samsung.com>
    CC: Maxime Ripard <mripard@kernel.org>
    CC: Chen-Yu Tsai <wens@csie.org>
    CC: Jernej Skrabec <jernej.skrabec@gmail.com>
    CC: Appana Durga Kedareswara rao <appana.durga.rao@xilinx.com>
    CC: Naga Sureshkumar Relli <naga.sureshkumar.relli@xilinx.com>
    CC: Michal Simek <michal.simek@xilinx.com>
    CC: Stephane Grosjean <s.grosjean@peak-system.com>
    Tested-by: Jimmy Assarsson <extja@kvaser.com> # kvaser
    Signed-off-by: Vincent Mailhol <mailhol.vincent@wanadoo.fr>
    Acked-by: Stefan Mätje <stefan.maetje@esd.eu> # esd_usb2
    Tested-by: Stefan Mätje <stefan.maetje@esd.eu> # esd_usb2
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit da2c624bab04df1bf42b71b5eeba8adc0809704d
Author: Joe Thornber <ejt@redhat.com>
Date:   Fri Dec 10 13:49:53 2021 +0000

    dm space map common: add bounds check to sm_ll_lookup_bitmap()

    [ Upstream commit cba23ac158db7f3cd48a923d6861bee2eb7a2978 ]

    Corrupted metadata could warrant returning error from sm_ll_lookup_bitmap().

    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 1b95c8dfe82e3cdf9a0b4af8a3bf6a12319c14e7
Author: Joe Thornber <ejt@redhat.com>
Date:   Fri Dec 10 13:44:13 2021 +0000

    dm btree: add a defensive bounds check to insert_at()

    [ Upstream commit 85bca3c05b6cca31625437eedf2060e846c4bbad ]

    Corrupt metadata could trigger an out of bounds write.

    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 96d936e9c05c7e8864ff9cc928a2d5c19f0ac3c3
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Mon Jan 3 09:36:21 2022 +0800

    mac80211: allow non-standard VHT MCS-10/11

    [ Upstream commit 04be6d337d37400ad5b3d5f27ca87645ee5a18a3 ]

    Some AP can possibly try non-standard VHT rate and mac80211 warns and drops
    packets, and leads low TCP throughput.

        Rate marked as a VHT rate but data is invalid: MCS: 10, NSS: 2
        WARNING: CPU: 1 PID: 7817 at net/mac80211/rx.c:4856 ieee80211_rx_list+0x223/0x2f0
[mac8021

    Since commit c27aa56a72b8 ("cfg80211: add VHT rate entries for MCS-10 and MCS-11")
```

```
    has added, mac80211 adds this support as well.

    After this patch, throughput is good and iw can get the bitrate:
        rx bitrate: 975.1 MBit/s VHT-MCS 10 80MHz short GI VHT-NSS 2
    or
        rx bitrate: 1083.3 MBit/s VHT-MCS 11 80MHz short GI VHT-NSS 2

    Buglink: https://bugzilla.suse.com/show_bug.cgi?id=1192891
    Reported-by: Goldwyn Rodrigues <rgoldwyn@suse.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Link: https://lore.kernel.org/r/20220103013623.17052-1-pkshih@realtek.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a094d44cb999c7407f07142f958a088ff1a28ec
Author: Florian Fainelli <f.fainelli@gmail.com>
Date:   Mon Jan 3 11:40:24 2022 -0800

    net: mdio: Demote probed message to debug print

    [ Upstream commit 7590fc6f80ac2cbf23e6b42b668bbeded070850b ]

    On systems with large numbers of MDIO bus/muxes the message indicating
    that a given MDIO bus has been successfully probed is repeated for as
    many buses we have, which can eat up substantial boot time for no
    reason, demote to a debug print.

    Reported-by: Maxime Bizon <mbizon@freebox.fr>
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Link: https://lore.kernel.org/r/20220103194024.2620-1-f.fainelli@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7e795cf927287cc054a291e8c652dbb30d7c6ca
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Fri Nov 5 16:45:35 2021 -0400

    btrfs: remove BUG_ON(!eie) in find_parent_nodes

    [ Upstream commit 9f05c09d6baef789726346397438cca4ec43c3ee ]

    If we're looking for leafs that point to a data extent we want to record
    the extent items that point at our bytenr.  At this point we have the
    reference and we know for a fact that this leaf should have a reference
    to our bytenr.  However if there's some sort of corruption we may not
    find any references to our leaf, and thus could end up with eie == NULL.
    Replace this BUG_ON() with an ASSERT() and then return -EUCLEAN for the
    mortals.

    Signed-off-by: Josef Bacik <josef@toxicpanda.com>
    Reviewed-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90fc21401f8e0dbefe3afbc9432500a09e3679a7
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Fri Nov 5 16:45:34 2021 -0400

    btrfs: remove BUG_ON() in find_parent_nodes()

    [ Upstream commit fcba0120edf88328524a4878d1d6f4ad39f2ec81 ]

    We search for an extent entry with .offset = -1, which shouldn't be a
    thing, but corruption happens.  Add an ASSERT() for the developers,
```

```
    return -EUCLEAN for mortals.

    Signed-off-by: Josef Bacik <josef@toxicpanda.com>
    Reviewed-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3894cf715ece7db93cf42e3727264b31ef5fc059
Author: Yevgeny Kliteynik <kliteyn@nvidia.com>
Date:   Tue Dec 14 12:56:18 2021 +0200

    net/mlx5: DR, Fix error flow in creating matcher

    [ Upstream commit 84dfac39c61fde04126e23723138128b50cd036f ]

    The error code of nic matcher init functions wasn't checked.
    This patch improves the matcher init function and fix error flow bug:
    the handling of match parameter is moved into a separate function
    and error flow is simplified.

    Signed-off-by: Yevgeny Kliteynik <kliteyn@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1d3076cf4315174033aa3a39dc5cd4b7b168780d
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Fri Dec 24 09:04:58 2021 +0800

    ACPI: CPPC: Check present CPUs for determining _CPC is valid

    [ Upstream commit 2aeca6bd02776d7f56a49a32be0dd184f204d888 ]

    As this is a static check, it should be based upon what is currently
    present on the system. This makes probeing more deterministic.

    While local APIC flags field (lapic_flags) of cpu core in MADT table is
    0, then the cpu core won't be enabled. In this case, _CPC won't be found
    in this core, and return back to _CPC invalid with walking through
    possible cpus (include disable cpus). This is not expected, so switch to
    check present CPUs instead.

    Reported-by: Jinzhou Su <Jinzhou.Su@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Huang Rui <ray.huang@amd.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8c8b8077441c7d2ca87fc59166d29f713ac16bc9
Author: Thomas Weißschuh <linux@weissschuh.net>
Date:   Wed Dec 22 22:20:14 2021 +0100

    ACPI: battery: Add the ThinkPad "Not Charging" quirk

    [ Upstream commit e96c1197aca628f7d2480a1cc3214912b40b3414 ]

    The EC/ACPI firmware on Lenovo ThinkPads used to report a status
    of "Unknown" when the battery is between the charge start and
    charge stop thresholds. On Windows, it reports "Not Charging"
    so the quirk has been added to also report correctly.

    Now the "status" attribute returns "Not Charging" when the
    battery on ThinkPads is not physicaly charging.

    Signed-off-by: Thomas Weißschuh <linux@weissschuh.net>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
```

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

commit 015ce2b573d2e94abdecfbf1e996c6307b835093
Author: Marina Nikolic <Marina.Nikolic@amd.com>
Date:   Tue Dec 14 20:57:53 2021 +0800

    amdgpu/pm: Make sysfs pm attributes as read-only for VFs

    [ Upstream commit 11c9cc95f818f0f187e9b579a7f136f532b42445 ]

    == Description ==
    Setting values of pm attributes through sysfs
    should not be allowed in SRIOV mode.
    These calls will not be processed by FW anyway,
    but error handling on sysfs level should be improved.

    == Changes ==
    This patch prohibits performing of all set commands
    in SRIOV mode on sysfs level.
    It offers better error handling as calls that are
    not allowed will not be propagated further.

    == Test ==
    Writing to any sysfs file in passthrough mode will succeed.
    Writing to any sysfs file in ONEVF mode will yield error:
    "calling process does not have sufficient permission to execute a command".

    Signed-off-by: Marina Nikolic <Marina.Nikolic@amd.com>
    Acked-by: Evan Quan <evan.quan@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 3f7fcff69710592fa2fd25a1f7d5069c1f2db074
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon Dec 6 18:48:06 2021 +0100

    mfd: intel_soc_pmic: Use CPU-id check instead of _HRV check to differentiate variants

    [ Upstream commit 5b78223f55a0f516a1639dbe11cd4324d4aaee20 ]

    The Intel Crystal Cove PMIC has 2 different variants, one for use with
    Bay Trail (BYT) SoCs and one for use with Cherry Trail (CHT) SoCs.

    So far we have been using an ACPI _HRV check to differentiate between
    the 2, but at least on the Microsoft Surface 3, which is a CHT device,
    the wrong _HRV value is reported by ACPI.

    So instead switch to a CPU-ID check which prevents us from relying on
    the possibly wrong ACPI _HRV value.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Reported-by: Tsuchiya Yuto <kitakar@gmail.com>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211206174806.197772-2-hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d0d0fdb733d1fe034cc50ddaa0c9bc42586cc23
Author: Zongmin Zhou <zhouzongmin@kylinos.cn>
Date:   Wed Dec 15 17:23:37 2021 +0800

    drm/amdgpu: fixup bad vram size on gmc v8

    [ Upstream commit 11544d77e3974924c5a9c8a8320b996a3e9b2f8b ]
```

```
    Some boards(like RX550) seem to have garbage in the upper
    16 bits of the vram size register.  Check for
    this and clamp the size properly.  Fixes
    boards reporting bogus amounts of vram.

    after add this patch,the maximum GPU VRAM size is 64GB,
    otherwise only 64GB vram size will be used.

    Signed-off-by: Zongmin Zhou<zhouzongmin@kylinos.cn>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8713ec3c70874b116f4d0b36670d18ea5568baca
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Thu Dec 16 13:57:44 2021 +0100

    mmc: mtk-sd: Use readl_poll_timeout instead of open-coded polling

    [ Upstream commit ffaea6ebfe9ce06ebb3a54811a47688f2b0893cd ]

    Replace all instances of open-coded while loops for polling registers
    with calls to readl_poll_timeout() and, while at it, also fix some
    possible infinite loop instances.

    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://lore.kernel.org/r/20211216125748.179602-1-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 761a564dc4761ec19ed2e5202fcf771b5fe3ab6b
Author: Kirill A. Shutemov <kirill.shutemov@linux.intel.com>
Date:   Wed Dec 22 17:33:51 2021 +0100

    ACPICA: Hardware: Do not flush CPU cache when entering S4 and S5

    [ Upstream commit 1d4e0b3abb168b2ee1eca99c527cffa1b80b6161 ]

    ACPICA commit 3dd7e1f3996456ef81bfe14cba29860e8d42949e

    According to ACPI 6.4, Section 16.2, the CPU cache flushing is
    required on entering to S1, S2, and S3, but the ACPICA code
    flushes the CPU cache regardless of the sleep state.

    Blind cache flush on entering S5 causes problems for TDX.

    Flushing happens with WBINVD that is not supported in the TDX
    environment.

    TDX only supports S5 and adjusting ACPICA code to conform to the
    spec more strictly fixes the issue.

    Link: https://github.com/acpica/acpica/commit/3dd7e1f3
    Signed-off-by: Kirill A. Shutemov <kirill.shutemov@linux.intel.com>
    [ rjw: Subject and changelog edits ]
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Bob Moore <robert.moore@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a54e17bd8d4015bce7249112ef40e86af01c1fa
Author: Sudeep Holla <sudeep.holla@arm.com>
Date:   Wed Dec 22 17:31:54 2021 +0100
```

```
    ACPICA: Fix wrong interpretation of PCC address

    [ Upstream commit 9a3b8655db1ada31c82189ae13f40eb25da48c35 ]

    ACPICA commit 41be6afacfdaec2dba3a5ed368736babc2a7aa5c

    With the PCC Opregion in the firmware and we are hitting below kernel crash:

    -->8
    Unable to handle kernel NULL pointer dereference at virtual address 0000000000000010
     Workqueue: pm pm_runtime_work
     pstate: 80000005 (Nzcv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
     pc : __memcpy+0x54/0x260
     lr : acpi_ex_write_data_to_field+0xb8/0x194
     Call trace:
      __memcpy+0x54/0x260
      acpi_ex_store_object_to_node+0xa4/0x1d4
      acpi_ex_store+0x44/0x164
      acpi_ex_opcode_1A_1T_1R+0x25c/0x508
      acpi_ds_exec_end_op+0x1b4/0x44c
      acpi_ps_parse_loop+0x3a8/0x614
      acpi_ps_parse_aml+0x90/0x2f4
      acpi_ps_execute_method+0x11c/0x19c
      acpi_ns_evaluate+0x1ec/0x2b0
      acpi_evaluate_object+0x170/0x2b0
      acpi_device_set_power+0x118/0x310
      acpi_dev_suspend+0xd4/0x180
      acpi_subsys_runtime_suspend+0x28/0x38
      __rpm_callback+0x74/0x328
      rpm_suspend+0x2d8/0x624
      pm_runtime_work+0xa4/0xb8
      process_one_work+0x194/0x25c
      worker_thread+0x260/0x49c
      kthread+0x14c/0x30c
      ret_from_fork+0x10/0x20
     Code: f9000006 f81f80a7 d65f03c0 361000c2 (b9400026)
     ---[ end trace 24d8a032fa77b68a ]---

    The reason for the crash is that the PCC channel index passed via region.address
    in acpi_ex_store_object_to_node is interpreted as the channel subtype
    incorrectly.

    Assuming the PCC op_region support is not used by any other type, let us
    remove the subtype check as the AML has no access to the subtype information.
    Once we remove it, the kernel crash disappears and correctly complains about
    missing PCC Opregion handler.

    ACPI Error: No handler for Region [PFRM] ((____ptrval____)) [PCC] (20210730/evregion-
130)
    ACPI Error: Region PCC (ID=10) has no handler (20210730/exfldio-261)
    ACPI Error: Aborting method \_SB.ETH0._PS3 due to previous error (AE_NOT_EXIST)
(20210730/psparse-531)

    Link: https://github.com/acpica/acpica/commit/41be6afa
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Bob Moore <robert.moore@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03d6ce317dc8e325e5236c26e9a6e9ba8c880082
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Wed Dec 22 17:31:05 2021 +0100

    ACPICA: Executer: Fix the REFCLASS_REFOF case in acpi_ex_opcode_1A_0T_1R()
```

```
    [ Upstream commit 24ea5f90ec9548044a6209685c5010edd66ffe8f ]

    ACPICA commit d984f12041392fa4156b52e2f7e5c5e7bc38ad9e

    If Operand[0] is a reference of the ACPI_REFCLASS_REFOF class,
    acpi_ex_opcode_1A_0T_1R () calls acpi_ns_get_attached_object () to
    obtain return_desc which may require additional resolution with
    the help of acpi_ex_read_data_from_field (). If the latter fails,
    the reference counter of the original return_desc is decremented
    which is incorrect, because acpi_ns_get_attached_object () does not
    increment the reference counter of the object returned by it.

    This issue may lead to premature deletion of the attached object
    while it is still attached and a use-after-free and crash in the
    host OS.  For example, this may happen when on evaluation of ref_of()
    a local region field where there is no registered handler for the
    given Operation Region.

    Fix it by making acpi_ex_opcode_1A_0T_1R () return Status right away
    after a acpi_ex_read_data_from_field () failure.

    Link: https://github.com/acpica/acpica/commit/d984f120
    Link: https://github.com/acpica/acpica/pull/685
    Reported-by: Lenny Szubowicz <lszubowi@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Bob Moore <robert.moore@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 445089528048fa929242487945fbf69aafdc475a
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Wed Dec 22 17:29:45 2021 +0100

    ACPICA: Utilities: Avoid deleting the same object twice in a row

    [ Upstream commit 1cdfe9e346b4c5509ffe19ccde880fd259d9f7a3 ]

    ACPICA commit c11af67d8f7e3d381068ce7771322f2b5324d687

    If original_count is 0 in acpi_ut_update_ref_count (),
    acpi_ut_delete_internal_obj () is invoked for the target object, which is
    incorrect, because that object has been deleted once already and the
    memory allocated to store it may have been reclaimed and allocated
    for a different purpose by the host OS.  Moreover, a confusing debug
    message following the "Reference Count is already zero, cannot
    decrement" warning is printed in that case.

    To fix this issue, make acpi_ut_update_ref_count () return after finding
    that original_count is 0 and printing the above warning.

    Link: https://github.com/acpica/acpica/commit/c11af67d
    Link: https://github.com/acpica/acpica/pull/652
    Reported-by: Mark Asselstine <mark.asselstine@windriver.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Bob Moore <robert.moore@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 648d75de59819cf48e179ed792b636cbb581574e
Author: Mark Langsdorf <mlangsdo@redhat.com>
Date:   Wed Dec 22 16:57:34 2021 +0100

    ACPICA: actypes.h: Expand the ACPI_ACCESS_ definitions

    [ Upstream commit f81bdeaf816142e0729eea0cc84c395ec9673151 ]
```

```
    ACPICA commit bc02c76d518135531483dfc276ed28b7ee632ce1

    The current ACPI_ACCESS_*_WIDTH defines do not provide a way to
    test that size is small enough to not cause an overflow when
    applied to a 32-bit integer.

    Rather than adding more magic numbers, add ACPI_ACCESS_*_SHIFT,
    ACPI_ACCESS_*_MAX, and ACPI_ACCESS_*_DEFAULT #defines and
    redefine ACPI_ACCESS_*_WIDTH in terms of the new #defines.

    This was inititally reported on Linux where a size of 102 in
    ACPI_ACCESS_BIT_WIDTH caused an overflow error in the SPCR
    initialization code.

    Link: https://github.com/acpica/acpica/commit/bc02c76d
    Signed-off-by: Mark Langsdorf <mlangsdo@redhat.com>
    Signed-off-by: Bob Moore <robert.moore@intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ee1f537cfb838fb1b4c080c3edf8989151f47db
Author: Kyeong Yoo <kyeong.yoo@alliedtelesis.co.nz>
Date:   Tue Jul 4 16:22:38 2017 +1200

    jffs2: GC deadlock reading a page that is used in jffs2_write_begin()

    [ Upstream commit aa39cc675799bc92da153af9a13d6f969c348e82 ]

    GC task can deadlock in read_cache_page() because it may attempt
    to release a page that is actually allocated by another task in
    jffs2_write_begin().
    The reason is that in jffs2_write_begin() there is a small window
    a cache page is allocated for use but not set Uptodate yet.

    This ends up with a deadlock between two tasks:
    1) A task (e.g. file copy)
       - jffs2_write_begin() locks a cache page
       - jffs2_write_end() tries to lock "alloc_sem" from
            jffs2_reserve_space() <-- STUCK
    2) GC task (jffs2_gcd_mtd3)
       - jffs2_garbage_collect_pass() locks "alloc_sem"
       - try to lock the same cache page in read_cache_page() <-- STUCK

    So to avoid this deadlock, hold "alloc_sem" in jffs2_write_begin()
    while reading data in a cache page.

    Signed-off-by: Kyeong Yoo <kyeong.yoo@alliedtelesis.co.nz>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e80561e50f7e1292ab0ca2abe57be1cbecdcfe92
Author: Lucas Stach <l.stach@pengutronix.de>
Date:   Wed Dec 22 01:17:28 2021 +0100

    drm/etnaviv: consider completed fence seqno in hang check

    [ Upstream commit cdd156955f946beaa5f3a00d8ccf90e5a197becc ]

    Some GPU heavy test programs manage to trigger the hangcheck quite often.
    If there are no other GPU users in the system and the test program
    exhibits a very regular structure in the commandstreams that are being
    submitted, we can end up with two distinct submits managing to trigger
    the hangcheck with the FE in a very similar address range. This leads
    the hangcheck to believe that the GPU is stuck, while in reality the GPU
```

```
    is already busy working on a different job. To avoid those spurious
    GPU resets, also remember and consider the last completed fence seqno
    in the hang check.

    Reported-by: Joerg Albert <joerg.albert@iav.de>
    Signed-off-by: Lucas Stach <l.stach@pengutronix.de>
    Reviewed-by: Christian Gmeiner <christian.gmeiner@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2aabcf5947192c743b74deb2146481142bfc7a57
Author: Antony Antony <antony.antony@secunet.com>
Date:   Wed Dec 22 14:11:18 2021 +0100

    xfrm: rate limit SA mapping change message to user space

    [ Upstream commit 4e484b3e969b52effd95c17f7a86f39208b2ccf4 ]

    Kernel generates mapping change message, XFRM_MSG_MAPPING,
    when a source port chage is detected on a input state with UDP
    encapsulation set.  Kernel generates a message for each IPsec packet
    with new source port.  For a high speed flow per packet mapping change
    message can be excessive, and can overload the user space listener.

    Introduce rate limiting for XFRM_MSG_MAPPING message to the user space.

    The rate limiting is configurable via netlink, when adding a new SA or
    updating it. Use the new attribute XFRMA_MTIMER_THRESH in seconds.

    v1->v2 change:
            update xfrm_sa_len()

    v2->v3 changes:
            use u32 insted unsigned long to reduce size of struct xfrm_state
            fix xfrm_ompat size Reported-by: kernel test robot <lkp@intel.com>
            accept XFRM_MSG_MAPPING only when XFRMA_ENCAP is present

    Co-developed-by: Thomas Egerer <thomas.egerer@secunet.com>
    Signed-off-by: Thomas Egerer <thomas.egerer@secunet.com>
    Signed-off-by: Antony Antony <antony.antony@secunet.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 118ff570ebafeb934f25e93849286172e3b5bccc
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Dec 22 12:22:01 2021 -0800

    Bluetooth: vhci: Set HCI_QUIRK_VALID_LE_STATES

    [ Upstream commit cfb4c313be670fd4bd09650216620fa4514cdb93 ]

    This set HCI_QUIRK_VALID_LE_STATES quirk which is required for the likes
    of experimental LE simultaneous roles.

    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 1878c20ba82fe614f367476ff94a136f0aa5a22a
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Fri Dec 17 20:06:08 2021 +0100

    cpufreq: intel_pstate: Update cpuinfo.max_freq on HWP_CAP changes

    [ Upstream commit dfeeedc1bf5772226bddf51ed3f853e5a6707bf1 ]
```

```
    With HWP enabled, when the turbo range of performance levels is
    disabled by the platform firmware, the CPU capacity is given by
    the "guaranteed performance" field in MSR_HWP_CAPABILITIES which
    is generally dynamic.  When it changes, the kernel receives an HWP
    notification interrupt handled by notify_hwp_interrupt().

    When the "guaranteed performance" value changes in the above
    configuration, the CPU performance scaling needs to be adjusted so
    as to use the new CPU capacity in computations, which means that
    the cpuinfo.max_freq value needs to be updated for that CPU.

    Accordingly, modify intel_pstate_notify_work() to read
    MSR_HWP_CAPABILITIES and update cpuinfo.max_freq to reflect the
    new configuration (this update can be carried out even if the
    configuration doesn't actually change, because it simply doesn't
    matter then and it takes less time to update it than to do extra
    checks to decide whether or not a change has really occurred).

    Reported-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Tested-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8458cf535fd33fadade5abe52b72f33b5139c3d0
Author: Tedd Ho-Jeong An <tedd.an@intel.com>
Date:   Mon Dec 13 23:32:14 2021 -0800

    Bluetooth: btintel: Add missing quirks and msft ext for legacy bootloader

    [ Upstream commit 3547a008c8962df2175db1e78b80f27e027ec549 ]

    This patch add missing HCI quirks and MSFT extension for legacy
    bootloader when it is running in the operational firmware.

    Signed-off-by: Tedd Ho-Jeong An <tedd.an@intel.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 b502bd83db4bc65c9dc5243413a5fb44d3bbf192
Author: Ben Greear <greearb@candelatech.com>
Date:   Thu Sep 3 12:52:54 2020 -0700

    ath11k: Fix napi related hang

    [ Upstream commit d943fdad7589653065be0e20aadc6dff37725ed4 ]

    Similar to the same bug in ath10k, a napi disable w/out it being enabled
    will hang forever.  I believe I saw this while trying rmmod after driver
    had some failure on startup.  Fix it by keeping state on whether napi is
    enabled or not.

    And, remove un-used napi pointer in ath11k driver base struct.

    Signed-off-by: Ben Greear <greearb@candelatech.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20200903195254.29379-1-greearb@candelatech.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8bd84be320dcbd489e22cd0469da6918f04c6897
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Sep 12 23:12:52 2021 -0700

    um: registers: Rename function names to avoid conflicts and build problems
```

```
    [ Upstream commit 077b7320942b64b0da182aefd83c374462a65535 ]

    The function names init_registers() and restore_registers() are used
    in several net/ethernet/ and gpu/drm/ drivers for other purposes (not
    calls to UML functions), so rename them.

    This fixes multiple build errors.

    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Jeff Dike <jdike@addtoit.com>
    Cc: Richard Weinberger <richard@nod.at>
    Cc: Anton Ivanov <anton.ivanov@cambridgegreys.com>
    Cc: linux-um@lists.infradead.org
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84da87f42daeb08e4dd72b6ca5d50043af02cb72
Author: Tetsuo Handa <penguin-kernel@i-love.sakura.ne.jp>
Date:   Fri Dec 17 23:51:25 2021 +0900

    block: check minor range in device_add_disk()

    [ Upstream commit e338924bd05d6e71574bc13e310c89e10e49a8a5 ]

    ioctl(fd, LOOP_CTL_ADD, 1048576) causes

      sysfs: cannot create duplicate filename '/dev/block/7:0'

    message because such request is treated as if ioctl(fd, LOOP_CTL_ADD, 0)
    due to MINORMASK == 1048575. Verify that all minor numbers for that device
    fit in the minor range.

    Reported-by: wangyangbo <wangyangbo@uniontech.com>
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/b1b19379-23ee-5379-0eb5-94bf5f79f1b4@i-love.sakura.ne.jp
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 02eda1806faafe2dd9a7ada1062381d9895446cc
Author: Hector Martin <marcan@marcan.st>
Date:   Thu Dec 16 01:10:44 2021 +0900

    mmc: sdhci-pci-gli: GL9755: Support for CD/WP inversion on OF platforms

    [ Upstream commit 189f1d9bc3a5ea3e442e119e4a5deda63da8c462 ]

    This is required on some Apple ARM64 laptops using this controller.
    As is typical on DT platforms, pull these quirks from the device tree
    using the standard mmc bindings.

    See Documentation/devicetree/bindings/mmc/mmc-controller.yaml

    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Link: https://lore.kernel.org/r/20211215161045.38843-2-marcan@marcan.st
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 74d27498bac549c5c3b36bf50a5b6b5042b9a08a
Author: Luca Coelho <luciano.coelho@intel.com>
Date:   Sun Dec 19 13:28:34 2021 +0200

    iwlwifi: pcie: make sure prph_info is set when treating wakeup IRQ
```

```
    [ Upstream commit 459fc0f2c6b0f6e280bfa0f230c100c9dfe3a199 ]

    In some rare cases when the HW is in a bad state, we may get this
    interrupt when prph_info is not set yet.  Then we will try to
    dereference it to check the sleep_notif element, which will cause an
    oops.

    Fix that by ignoring the interrupt if prph_info is not set yet.

    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219132536.0537aa562313.I183bb336345b9b3da196ba9e596a6
f189fbcbd09@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68a0f06e607f75bbd5f4824ec99489e13128a80f
Author: Avraham Stern <avraham.stern@intel.com>
Date:   Sun Dec 19 13:28:31 2021 +0200

    iwlwifi: mvm: fix AUX ROC removal

    [ Upstream commit f0337cb48f3bf5f0bbccc985d8a0a8c4aa4934b7 ]

    When IWL_UCODE_TLV_CAPA_SESSION_PROT_CMD is set, removing a time event
    always tries to cancel session protection. However, AUX ROC does
    not use session protection so it was never removed. As a result,
    if the driver tries to allocate another AUX ROC event right after
    cancelling the first one, it will fail with a warning.
    In addition, the time event data passed to iwl_mvm_remove_aux_roc_te()
    is incorrect. Fix it.

    Signed-off-by: Avraham Stern <avraham.stern@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219132536.915e1f69f062.Id837e917f1c2beaca7c1eb33333d6
22548918a76@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7fa7a30fa86db624b7850434e6da8d3c07ea2891
Author: Ilan Peer <ilan.peer@intel.com>
Date:   Sun Dec 19 12:18:16 2021 +0200

    iwlwifi: mvm: Fix calculation of frame length

    [ Upstream commit 40a0b38d7a7f91a6027287e0df54f5f547e8d27e ]

    The RADA might include in the Rx frame the MIC and CRC bytes.
    These bytes should be removed for non monitor interfaces and
    should not be passed to mac80211.

    Fix the Rx processing to remove the extra bytes on non monitor
    cases.

    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.20211219121514.098be12c801e.I1d81733d8a75b84c3b20eb6e0d14a
b3405ca6a86@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b54dc603bfc691a15d261c8d33c074b44a582bc2
Author: Johannes Berg <johannes.berg@intel.com>
```

```
Date:   Fri Dec 10 11:12:45 2021 +0200

    iwlwifi: remove module loading failure message

    [ Upstream commit 6518f83ffa51131daaf439b66094f684da3fb0ae ]

    When CONFIG_DEBUG_TEST_DRIVER_REMOVE is set, iwlwifi crashes
    when the opmode module cannot be loaded, due to completing
    the completion before using drv->dev, which can then already
    be freed.

    Fix this by removing the (fairly useless) message. Moving the
    completion later causes a deadlock instead, so that's not an
    option.

    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link: https://lore.kernel.org/r/20211210091245.289008-2-luca@coelho.fi
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6b5ad4bd0d78fef6bbe0ecdf96e09237c9c52cc1
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Fri Dec 10 11:12:42 2021 +0200

    iwlwifi: fix leaks/bad data after failed firmware load

    [ Upstream commit ab07506b0454bea606095951e19e72c282bfbb42 ]

    If firmware load fails after having loaded some parts of the
    firmware, e.g. the IML image, then this would leak. For the
    host command list we'd end up running into a WARN on the next
    attempt to load another firmware image.

    Fix this by calling iwl_dealloc_ucode() on failures, and make
    that also clear the data so we start fresh on the next round.

    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.20211210110539.1f742f0eb58a.I1315f22f6aa632d94ae2069f85e1b
ca5e734dce0@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f62db28fbfaff6a40d150aeb56d9c4f22c92328b
Author: Luca Coelho <luciano.coelho@intel.com>
Date:   Fri Dec 10 09:06:20 2021 +0200

    iwlwifi: recognize missing PNVM data and then log filename

    [ Upstream commit 1db385c668d3ccb04ccdb5f0f190fd17b2db29c6 ]

    We can detect that a FW SYSASSERT is due to missing PNVM data by
    checking the assertion code.  When this happens, it's is useful for
    the user if we print the filename where the driver is looking for the
    data.

    Add the PNVM missing assertion code to the dump list and print out the
    name of the file we're looking for when this happens.

    Reported-by: Sam Edwards <CFSworks@gmail.com>
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211210090244.1d8725b7518a.I0c36617a7282bd445cda484d97ac4
```

```
a83022706ee@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1a88b7960a3c3b71c1f5071a299a132e41ade25
Author: Changcheng Deng <deng.changcheng@zte.com.cn>
Date:   Thu Nov 25 01:43:11 2021 +0000

    PM: AVS: qcom-cpr: Use div64_ul instead of do_div

    [ Upstream commit 92c550f9ffd2884bb5def52b5c0485a35e452784 ]

    do_div() does a 64-by-32 division. Here the divisor is an unsigned long
    which on some platforms is 64 bit wide. So use div64_ul instead of do_div
    to avoid a possible truncation.

    Reported-by: Zeal Robot <zealci@zte.com.cn>
    Signed-off-by: Changcheng Deng <deng.changcheng@zte.com.cn>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211125014311.45942-1-deng.changcheng@zte.com.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4b986bd6004cf07c28d6335b685f795cb94a81e1
Author: Po-Hao Huang <phhuang@realtek.com>
Date:   Fri Dec 17 09:27:08 2021 +0800

    rtw88: 8822c: update rx settings to prevent potential hw deadlock

    [ Upstream commit c1afb26727d9e507d3e17a9890e7aaf7fc85cd55 ]

    These settings enables mac to detect and recover when rx fifo
    circuit deadlock occurs. Previous version missed this, so we fix it.

    Signed-off-by: Po-Hao Huang <phhuang@realtek.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20211217012708.8623-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1081a9ac20c1721c080a582bce6307c4fe6e42e1
Author: Zekun Shen <bruceshenzk@gmail.com>
Date:   Thu Oct 28 18:21:42 2021 -0400

    ath9k: Fix out-of-bound memcpy in ath9k_hif_usb_rx_stream

    [ Upstream commit 6ce708f54cc8d73beca213cec66ede5ce100a781 ]

    Large pkt_len can lead to out-out-bound memcpy. Current
    ath9k_hif_usb_rx_stream allows combining the content of two urb
    inputs to one pkt. The first input can indicate the size of the
    pkt. Any remaining size is saved in hif_dev->rx_remain_len.
    While processing the next input, memcpy is used with rx_remain_len.

    4-byte pkt_len can go up to 0xffff, while a single input is 0x4000
    maximum in size (MAX_RX_BUF_SIZE). Thus, the patch adds a check for
    pkt_len which must not exceed 2 * MAX_RX_BUG_SIZE.

    BUG: KASAN: slab-out-of-bounds in ath9k_hif_usb_rx_cb+0x490/0xed7 [ath9k_htc]
    Read of size 46393 at addr ffff888018798000 by task kworker/0:1/23

    CPU: 0 PID: 23 Comm: kworker/0:1 Not tainted 5.6.0 #63
    Hardware name: QEMU Standard PC (i440FX + PIIX, 1996),
    BIOS rel-1.10.2-0-g5f4c7b1-prebuilt.qemu-project.org 04/01/2014
    Workqueue: events request_firmware_work_func
    Call Trace:
```

```
       <IRQ>
       dump_stack+0x76/0xa0
       print_address_description.constprop.0+0x16/0x200
       ? ath9k_hif_usb_rx_cb+0x490/0xed7 [ath9k_htc]
       ? ath9k_hif_usb_rx_cb+0x490/0xed7 [ath9k_htc]
       __kasan_report.cold+0x37/0x7c
       ? ath9k_hif_usb_rx_cb+0x490/0xed7 [ath9k_htc]
       kasan_report+0xe/0x20
       check_memory_region+0x15a/0x1d0
       memcpy+0x20/0x50
       ath9k_hif_usb_rx_cb+0x490/0xed7 [ath9k_htc]
       ? hif_usb_mgmt_cb+0x2d9/0x2d9 [ath9k_htc]
       ? _raw_spin_lock_irqsave+0x7b/0xd0
       ? _raw_spin_trylock_bh+0x120/0x120
       ? __usb_unanchor_urb+0x12f/0x210
       __usb_hcd_giveback_urb+0x1e4/0x380
       usb_giveback_urb_bh+0x241/0x4f0
       ? __hrtimer_run_queues+0x316/0x740
       ? __usb_hcd_giveback_urb+0x380/0x380
       tasklet_action_common.isra.0+0x135/0x330
       __do_softirq+0x18c/0x634
       irq_exit+0x114/0x140
       smp_apic_timer_interrupt+0xde/0x380
       apic_timer_interrupt+0xf/0x20

    I found the bug using a custome USBFuzz port. It's a research work
    to fuzz USB stack/drivers. I modified it to fuzz ath9k driver only,
    providing hand-crafted usb descriptors to QEMU.

    After fixing the value of pkt_tag to ATH_USB_RX_STREAM_MODE_TAG in QEMU
    emulation, I found the KASAN report. The bug is triggerable whenever
    pkt_len is above two MAX_RX_BUG_SIZE. I used the same input that crashes
    to test the driver works when applying the patch.

    Signed-off-by: Zekun Shen <bruceshenzk@gmail.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/YXsidrRuK6zBJicZ@10-18-43-117.dynapool.wireless.nyu.edu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7bbc1a50a7963f14048f0e54b0b73159f86d4ea3
Author: Tetsuo Handa <penguin-kernel@i-love.sakura.ne.jp>
Date:   Tue Sep 21 22:06:26 2021 +0900

    ath9k_htc: fix NULL pointer dereference at ath9k_htc_tx_get_packet()

    [ Upstream commit 8b3046abc99eefe11438090bcc4ec3a3994b55d0 ]

    syzbot is reporting lockdep warning at ath9k_wmi_event_tasklet() followed
    by kernel panic at get_htc_epid_queue() from ath9k_htc_tx_get_packet() from
    ath9k_htc_txstatus() [1], for ath9k_wmi_event_tasklet(WMI_TXSTATUS_EVENTID)
    depends on spin_lock_init() from ath9k_init_priv() being already completed.

    Since ath9k_wmi_event_tasklet() is set by ath9k_init_wmi() from
    ath9k_htc_probe_device(), it is possible that ath9k_wmi_event_tasklet() is
    called via tasklet interrupt before spin_lock_init() from ath9k_init_priv()
     from ath9k_init_device() from ath9k_htc_probe_device() is called.

    Let's hold ath9k_wmi_event_tasklet(WMI_TXSTATUS_EVENTID) no-op until
    ath9k_tx_init() completes.

    Link: https://syzkaller.appspot.com/bug?extid=31d54c60c5b254d6f75b [1]
    Reported-by: syzbot <syzbot+31d54c60c5b254d6f75b@syzkaller.appspotmail.com>
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Tested-by: syzbot <syzbot+31d54c60c5b254d6f75b@syzkaller.appspotmail.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
```

```
    Link: https://lore.kernel.org/r/77b76ac8-2bee-6444-d26c-8c30858b8daa@i-love.sakura.ne.jp
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b9913a8bd7a83da10bb72d57d4d52a165fbf408
Author: Tetsuo Handa <penguin-kernel@i-love.sakura.ne.jp>
Date:   Tue Sep 21 22:06:23 2021 +0900

    ath9k_htc: fix NULL pointer dereference at ath9k_htc_rxep()

    [ Upstream commit b0ec7e55fce65f125bd1d7f02e2dc4de62abee34 ]

    syzbot is reporting lockdep warning followed by kernel panic at
    ath9k_htc_rxep() [1], for ath9k_htc_rxep() depends on ath9k_rx_init()
    being already completed.

    Since ath9k_htc_rxep() is set by ath9k_htc_connect_svc(WMI_BEACON_SVC)
     from ath9k_init_htc_services(), it is possible that ath9k_htc_rxep() is
    called via timer interrupt before ath9k_rx_init() from ath9k_init_device()
    is called.

    Since we can't call ath9k_init_device() before ath9k_init_htc_services(),
    let's hold ath9k_htc_rxep() no-op until ath9k_rx_init() completes.

    Link: https://syzkaller.appspot.com/bug?extid=4d2d56175b934b9a7bf9 [1]
    Reported-by: syzbot <syzbot+4d2d56175b934b9a7bf9@syzkaller.appspotmail.com>
    Signed-off-by: Tetsuo Handa <penguin-kernel@I-love.SAKURA.ne.jp>
    Tested-by: syzbot <syzbot+4d2d56175b934b9a7bf9@syzkaller.appspotmail.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/2b88f416-b2cb-7a18-d688-951e6dc3fe92@i-love.sakura.ne.jp
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f605eb623902cf0f2cac635e4fbb13ce3174888
Author: Deren Wu <deren.wu@mediatek.com>
Date:   Wed Dec 8 22:59:55 2021 +0800

    mt76: mt7921: fix network buffer leak by txs missing

    [ Upstream commit e0bf699ad8e52330d848144a9624a067ad588bd6 ]

    TXS in mt7921 may be forwared to tx_done event. Should try to catch
    TXS information in tx_done event as well.

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

commit 7e406591b77cc73135b8a8d09559fa128d30c4fa
Author: Felix Fietkau <nbd@nbd.name>
Date:   Thu Nov 25 10:42:12 2021 +0100

    mt76: mt7615: improve wmm index allocation

    [ Upstream commit 70fb028707c8871ef9e56b3ffa3d895e14956cc4 ]

    Typically all AP interfaces on a PHY will share the same WMM settings, while
    sta/mesh interfaces will usually inherit the settings from a remote device.
    In order minimize the likelihood of conflicting WMM settings, make all AP
    interfaces share one slot, and all non-AP interfaces another one.

    This also fixes running multiple AP interfaces on MT7613, which only has 3
    WMM slots.

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

```
commit 97cd90029d3dce8f851082316b901d0197a20835
Author: Xing Song <xing.song@mediatek.com>
Date:   Fri Nov 19 14:37:06 2021 +0800

    mt76: do not pass the received frame with decryption error

    [ Upstream commit dd28dea52ad9376d2b243a8981726646e1f60b1a ]

    MAC80211 doesn't care any decryption error in 802.3 path, so received
    frame will be dropped if HW tell us that the cipher configuration is not
    matched as well as the header has been translated to 802.3. This case only
    appears when IEEE80211_FCTL_PROTECTED is 0 and cipher suit is not none in
    the corresponding HW entry.

    The received frame is only reported to monitor interface if HW decryption
    block tell us there is ICV error or CCMP/BIP/WPI MIC error. Note in this
    case the reported frame is decrypted 802.11 frame and the payload may be
    malformed due to mismatched key.

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

commit 2f4d6700e416a5cec32e371af6b7b69e4dde4c20
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Nov 18 13:49:41 2021 +0100

    mt76: connac: fix a theoretical NULL pointer dereference in mt76_connac_get_phy_mode

    [ Upstream commit c9dbeac4988f6d4c4211b3747ec9c9c75ea87967 ]

    Even if it is not a real bug since mt76_connac_get_phy_mode runs just
    for mt7921 where only STA is supported, fix a theoretical NULL pointer
    dereference if new added modes do not support HE

    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 3a1c6734e5f3263a0956e66699116514fc7c3a77
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Fri Nov 12 03:45:31 2021 +0800

    mt76: mt7915: fix SMPS operation fail

    [ Upstream commit 8f05835425ce3f669e4b6d7c2c39a9aa22e1506c ]

    TGn fails sending SM power save mode action frame to the AP to switch
    from dynamic SMPS mode to static mode.

    Reported-by: Fang Zhao <fang.zhao@mediatek.com>
    Signed-off-by: Fang Zhao <fang.zhao@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 ab875e2bcc88249f5cf722aa0c4bb1a828daa256
Author: Peter Chiu <chui-hao.chiu@mediatek.com>
Date:   Mon Nov 1 10:01:13 2021 +0800

    mt76: mt7615: fix possible deadlock while mt7615_register_ext_phy()

    [ Upstream commit 8c55516de3f9b76b9d9444e7890682ec2efc809f ]

    ieee80211_register_hw() is called with rtnl_lock held, and this could be
```

```
    caused lockdep from a work item that's on a workqueue that is flushed
    with the rtnl held.

    Move mt7615_register_ext_phy() outside the init_work().

    Signed-off-by: Peter Chiu <chui-hao.chiu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a7992e29e569084344c8901c8dc05068e14ef5f
Author: Kai-Heng Feng <kai.heng.feng@canonical.com>
Date:   Wed Dec 15 20:01:06 2021 +0800

    usb: hub: Add delay for SuperSpeed hub resume to let links transit to U0

    [ Upstream commit 00558586382891540c59c9febc671062425a6e47 ]

    When a new USB device gets plugged to nested hubs, the affected hub,
    which connects to usb 2-1.4-port2, doesn't report there's any change,
    hence the nested hubs go back to runtime suspend like nothing happened:
    [  281.032951] usb usb2: usb wakeup-resume
    [  281.032959] usb usb2: usb auto-resume
    [  281.032974] hub 2-0:1.0: hub_resume
    [  281.033011] usb usb2-port1: status 0263 change 0000
    [  281.033077] hub 2-0:1.0: state 7 ports 4 chg 0000 evt 0000
    [  281.049797] usb 2-1: usb wakeup-resume
    [  281.069800] usb 2-1: Waited 0ms for CONNECT
    [  281.069810] usb 2-1: finish resume
    [  281.070026] hub 2-1:1.0: hub_resume
    [  281.070250] usb 2-1-port4: status 0203 change 0000
    [  281.070272] usb usb2-port1: resume, status 0
    [  281.070282] hub 2-1:1.0: state 7 ports 4 chg 0010 evt 0000
    [  281.089813] usb 2-1.4: usb wakeup-resume
    [  281.109792] usb 2-1.4: Waited 0ms for CONNECT
    [  281.109801] usb 2-1.4: finish resume
    [  281.109991] hub 2-1.4:1.0: hub_resume
    [  281.110147] usb 2-1.4-port2: status 0263 change 0000
    [  281.110234] usb 2-1-port4: resume, status 0
    [  281.110239] usb 2-1-port4: status 0203, change 0000, 10.0 Gb/s
    [  281.110266] hub 2-1.4:1.0: state 7 ports 4 chg 0000 evt 0000
    [  281.110426] hub 2-1.4:1.0: hub_suspend
    [  281.110565] usb 2-1.4: usb auto-suspend, wakeup 1
    [  281.130998] hub 2-1:1.0: hub_suspend
    [  281.137788] usb 2-1: usb auto-suspend, wakeup 1
    [  281.142935] hub 2-0:1.0: state 7 ports 4 chg 0000 evt 0000
    [  281.177828] usb 2-1: usb wakeup-resume
    [  281.197839] usb 2-1: Waited 0ms for CONNECT
    [  281.197850] usb 2-1: finish resume
    [  281.197984] hub 2-1:1.0: hub_resume
    [  281.198203] usb 2-1-port4: status 0203 change 0000
    [  281.198228] usb usb2-port1: resume, status 0
    [  281.198237] hub 2-1:1.0: state 7 ports 4 chg 0010 evt 0000
    [  281.217835] usb 2-1.4: usb wakeup-resume
    [  281.237834] usb 2-1.4: Waited 0ms for CONNECT
    [  281.237845] usb 2-1.4: finish resume
    [  281.237990] hub 2-1.4:1.0: hub_resume
    [  281.238067] usb 2-1.4-port2: status 0263 change 0000
    [  281.238148] usb 2-1-port4: resume, status 0
    [  281.238152] usb 2-1-port4: status 0203, change 0000, 10.0 Gb/s
    [  281.238166] hub 2-1.4:1.0: state 7 ports 4 chg 0000 evt 0000
    [  281.238385] hub 2-1.4:1.0: hub_suspend
    [  281.238523] usb 2-1.4: usb auto-suspend, wakeup 1
    [  281.258076] hub 2-1:1.0: hub_suspend
    [  281.265744] usb 2-1: usb auto-suspend, wakeup 1
    [  281.285976] hub 2-0:1.0: hub_suspend
```

```
    [  281.285988] usb usb2: bus auto-suspend, wakeup 1

    USB 3.2 spec, 9.2.5.4 "Changing Function Suspend State" says that "If
    the link is in a non-U0 state, then the device must transition the link
    to U0 prior to sending the remote wake message", but the hub only
    transits the link to U0 after signaling remote wakeup.

    So be more forgiving and use a 20ms delay to let the link transit to U0
    for remote wakeup.

    Suggested-by: Alan Stern <stern@rowland.harvard.edu>
    Acked-by: Alan Stern <stern@rowland.harvard.edu>
    Signed-off-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Link: https://lore.kernel.org/r/20211215120108.336597-1-kai.heng.feng@canonical.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f370288ea3044cd40d795f85226a6820b6b01ada
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Thu Dec 16 20:32:15 2021 +0100

    cpufreq: Fix initialization of min and max frequency QoS requests

    [ Upstream commit 521223d8b3ec078f670c7c35a1a04b1b2af07966 ]

    The min and max frequency QoS requests in the cpufreq core are
    initialized to whatever the current min and max frequency values are
    at the init time, but if any of these values change later (for
    example, cpuinfo.max_freq is updated by the driver), these initial
    request values will be limiting the CPU frequency unnecessarily
    unless they are changed by user space via sysfs.

    To address this, initialize min_freq_req and max_freq_req to
    FREQ_QOS_MIN_DEFAULT_VALUE and FREQ_QOS_MAX_DEFAULT_VALUE,
    respectively, so they don't really limit anything until user
    space updates them.

    Reported-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Tested-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 49dd63a6ddd396c8b9c9f27d03358771ac7145dc
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Fri Dec 10 17:10:13 2021 +0100

    PM: runtime: Add safety net to supplier device release

    [ Upstream commit d1579e61192e0e686faa4208500ef4c3b529b16c ]

    Because refcount_dec_not_one() returns true if the target refcount
    becomes saturated, it is generally unsafe to use its return value as
    a loop termination condition, but that is what happens when a device
    link's supplier device is released during runtime PM suspend
    operations and on device link removal.

    To address this, introduce pm_runtime_release_supplier() to be used
    in the above cases which will check the supplier device's runtime
    PM usage counter in addition to the refcount_dec_not_one() return
    value, so the loop can be terminated in case the rpm_active refcount
    value becomes invalid, and update the code in question to use it as
    appropriate.

    This change is not expected to have any visible functional impact.
```

```
    Reported-by: Peter Zijlstra <peterz@infradead.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Acked-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Acked-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 024ccff9af9cc59eb9fcd011013c7aca64c9d1c2
Author: Yang Shen <shenyang39@huawei.com>
Date:   Sat Dec 11 16:56:55 2021 +0800

    crypto: hisilicon/qm - fix deadlock for remove driver

    [ Upstream commit fc6c01f0cd10b89c4b01dd2940e0b0cda1bd82fb ]

    When remove the driver and executing the task occur at the same time,
    the following deadlock will be triggered:

    Chain exists of:
        sva_lock --> uacce_mutex --> &qm->qps_lock
        Possible unsafe locking scenario:
                    CPU0                    CPU1
                    ----                    ----
            lock(&qm->qps_lock);
                                            lock(uacce_mutex);
                                            lock(&qm->qps_lock);
            lock(sva_lock);

    And the lock 'qps_lock' is used to protect qp. Therefore, it's reasonable
    cycle is to continue until the qp memory is released. So move the release
    lock infront of 'uacce_remove'.

    Signed-off-by: Yang Shen <shenyang39@huawei.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d9b5afc35bcd451775621c1a3432bb6dca94f47b
Author: Weili Qian <qianweili@huawei.com>
Date:   Sat Dec 11 16:17:19 2021 +0800

    crypto: hisilicon/hpre - fix memory leak in hpre_curve25519_src_init()

    [ Upstream commit 51fa916b81e5f406a74f14a31a3a228c3cc060ad ]

    hpre_curve25519_src_init() allocates memory for 'ptr' before calling
    memcmp(). If memcmp() returns 0, the function will return '-EINVAL'
    without freeing memory.

    Signed-off-by: Weili Qian <qianweili@huawei.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit decce23b3623bd4f6c9b86ff5dd8c8d9cd544c69
Author: Peter Gonda <pgonda@google.com>
Date:   Tue Dec 7 15:33:03 2021 -0800

    crypto: ccp - Move SEV_INIT retry for corrupted data

    [ Upstream commit e423b9d75e779d921e6adf5ac3d0b59400d6ba7e ]

    Move the data corrupted retry of SEV_INIT into the
    __sev_platform_init_locked() function. This is for upcoming INIT_EX
    support as well as helping direct callers of
    __sev_platform_init_locked() which currently do not support the
    retry.
```

```
    Signed-off-by: Peter Gonda <pgonda@google.com>
    Reviewed-by: Marc Orr <marcorr@google.com>
    Acked-by: David Rientjes <rientjes@google.com>
    Acked-by: Tom Lendacky <thomas.lendacky@amd.com>
    Acked-by: Brijesh Singh <brijesh.singh@amd.com>
    Cc: Tom Lendacky <thomas.lendacky@amd.com>
    Cc: Brijesh Singh <brijesh.singh@amd.com>
    Cc: Marc Orr <marcorr@google.com>
    Cc: Joerg Roedel <jroedel@suse.de>
    Cc: Herbert Xu <herbert@gondor.apana.org.au>
    Cc: David Rientjes <rientjes@google.com>
    Cc: John Allen <john.allen@amd.com>
    Cc: "David S. Miller" <davem@davemloft.net>
    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Cc: linux-crypto@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bd478311ef7d195f4621967e23601162bf0794dd
Author: Thierry Reding <treding@nvidia.com>
Date:   Tue Dec 7 14:28:29 2021 +0100

    arm64: tegra: Adjust length of CCPLEX cluster MMIO region

    [ Upstream commit 2b14cbd643feea5fc17c6e8bead4e71088c69acd ]

    The Tegra186 CCPLEX cluster register region is 4 MiB is length, not 4
    MiB - 1. This was likely presumed to be the "limit" rather than length.
    Fix it up.

    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba202cf0d40084ca98c98ddbeeb7523836b12743
Author: Biwen Li <biwen.li@nxp.com>
Date:   Tue Dec 14 03:32:38 2021 -0600

    arm64: dts: ls1028a-qds: move rtc node to the correct i2c bus

    [ Upstream commit cbe9d948eadfe352ad45495a7cc5bf20a1b29d90 ]

    The i2c rtc is on i2c2 bus not i2c1 bus, so fix it in dts.

    Signed-off-by: Biwen Li <biwen.li@nxp.com>
    Signed-off-by: Li Yang <leoyang.lil@nxp.com>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9748204c9d6e99ebfd309420cc2d21402448a374
Author: Paul Moore <paul@paul-moore.com>
Date:   Mon Dec 13 15:45:20 2021 -0500

    audit: ensure userspace is penalized the same as the kernel when under pressure

    [ Upstream commit 8f110f530635af44fff1f4ee100ecef0bac62510 ]

    Due to the audit control mutex necessary for serializing audit
    userspace messages we haven't been able to block/penalize userspace
    processes that attempt to send audit records while the system is
    under audit pressure.  The result is that privileged userspace
    applications have a priority boost with respect to audit as they are
    not bound by the same audit queue throttling as the other tasks on
    the system.
```

```
    This patch attempts to restore some balance to the system when under
    audit pressure by blocking these privileged userspace tasks after
    they have finished their audit processing, and dropped the audit
    control mutex, but before they return to userspace.

    Reported-by: Gaosheng Cui <cuigaosheng1@huawei.com>
    Tested-by: Gaosheng Cui <cuigaosheng1@huawei.com>
    Reviewed-by: Richard Guy Briggs <rgb@redhat.com>
    Signed-off-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3038964acbda39ac7e29e305989fb10b3aaefdbc
Author: Jingwen Chen <Jingwen.Chen2@amd.com>
Date:   Tue Dec 14 11:50:39 2021 +0800

    drm/amd/amdgpu: fix gmc bo pin count leak in SRIOV

    [ Upstream commit 948e7ce01413b71395723aaf846015062aea3a43 ]

    [Why]
    gmc bo will be pinned during loading amdgpu and reset in SRIOV while
    only unpinned in unload amdgpu

    [How]
    add amdgpu_in_reset and sriov judgement to skip pin bo

    v2: fix wrong judgement

    Signed-off-by: Jingwen Chen <Jingwen.Chen2@amd.com>
    Reviewed-by: Horace Chen <horace.chen@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc0a5bfbf941dba7b266d6004f83d7cf89364a8e
Author: Jingwen Chen <Jingwen.Chen2@amd.com>
Date:   Tue Dec 14 11:31:16 2021 +0800

    drm/amd/amdgpu: fix psp tmr bo pin count leak in SRIOV

    [ Upstream commit 85dfc1d692c9434c37842e610be37cd4ae4e0081 ]

    [Why]
    psp tmr bo will be pinned during loading amdgpu and reset in SRIOV while
    only unpinned in unload amdgpu

    [How]
    add amdgpu_in_reset and sriov judgement to skip pin bo

    v2: fix wrong judgement

    Signed-off-by: Jingwen Chen <Jingwen.Chen2@amd.com>
    Reviewed-by: Horace Chen <horace.chen@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1ae48c70e59f56ecae5c88d32d58f5038e85ea6
Author: H. Nikolaus Schaller <hns@goldelico.com>
Date:   Wed Nov 10 18:17:11 2021 +0100

    mmc: omap_hsmmc: Revert special init for wl1251

    [ Upstream commit dfb654f1885f05baf506cdfdbc3f7efa1d847d54 ]

    Due to recent changes to the mmc core, card quirks can be applied based
    upon a compatible string in child OF node. The quirk needed for wl1251
```

```
    (SDIO card) is managed in the core, therefore there is no longer any reason
    to deal with this in omap_hsmmc too, so let's remove it.

    Signed-off-by: H. Nikolaus Schaller <hns@goldelico.com>
    Link:
https://lore.kernel.org/r/77d313b97d1e18b0eb7ed2d88d718d960f329bb0.1636564631.git.hns@goldel
ico.com
    [Ulf: Re-wrote the commit message to make it more clear]
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b42be87cb5bda29c5d6ea7ba19adb0ff53e7250f
Author: Ulf Hansson <ulf.hansson@linaro.org>
Date:   Wed Nov 10 18:17:09 2021 +0100

    mmc: core: Fixup storing of OCR for MMC_QUIRK_NONSTD_SDIO

    [ Upstream commit 8c3e5b74b9e2146f564905e50ca716591c76d4f1 ]

    The mmc core takes a specific path to support initializing of a
    non-standard SDIO card. This is triggered by looking for the card-quirk,
    MMC_QUIRK_NONSTD_SDIO.

    In mmc_sdio_init_card() this gets rather messy, as it causes the code to
    bail out earlier, compared to the usual path. This leads to that the OCR
    doesn't get saved properly in card->ocr. Fortunately, only omap_hsmmc has
    been using the MMC_QUIRK_NONSTD_SDIO and is dealing with the issue, by
    assigning a hardcoded value (0x80) to card->ocr from an ->init_card() ops.

    To make the behaviour consistent, let's instead rely on the core to save
    the OCR in card->ocr during initialization.

    Reported-by: H. Nikolaus Schaller <hns@goldelico.com>
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: H. Nikolaus Schaller <hns@goldelico.com>
    Link:
https://lore.kernel.org/r/e7936cff7fc24d187ef2680d3b4edb0ade58f293.1636564631.git.hns@goldel
ico.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 287fff87a6b81f365b148bb8dcb4eea3f089f383
Author: Biju Das <biju.das.jz@bp.renesas.com>
Date:   Wed Nov 3 13:26:46 2021 +0100

    mmc: tmio: reinit card irqs in reset routine

    [ Upstream commit e315b1f3a170f368da5618f8a598e68880302ed1 ]

    Refactor the code so that card detect irqs are always reenabled after a
    reset. This avoids doing it manually all over the code or forgetting to
    do this in the future.

    Reported-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Signed-off-by: Biju Das <biju.das.jz@bp.renesas.com>
    [wsa: added a comment when 'native_hotplug' has to be set]
    Signed-off-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Link: https://lore.kernel.org/r/20211103122646.64422-1-wsa+renesas@sang-engineering.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05e0f3be3c7dcd44ca17886b4537a67261bfb1c5
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Fri Dec 3 16:40:30 2021 +0100
```

```
    media: saa7146: hexium_gemini: Fix a NULL pointer dereference in hexium_attach()

    [ Upstream commit 3af86b046933ba513d08399dba0d4d8b50d607d0 ]

    In hexium_attach(dev, info), saa7146_vv_init() is called to allocate
    a new memory for dev->vv_data. saa7146_vv_release() will be called on
    failure of saa7146_register_device(). There is a dereference of
    dev->vv_data in saa7146_vv_release(), which could lead to a NULL
    pointer dereference on failure of saa7146_vv_init().

    Fix this bug by adding a check of saa7146_vv_init().

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_VIDEO_HEXIUM_GEMINI=m show no new warnings,
    and our static analyzer no longer warns about this code.

    Link: https://lore.kernel.org/linux-media/20211203154030.111210-1-zhou1615@umn.edu
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec27b051b13195115ff2e25a0b505029e784ea12
Author: Mikhail Rudenko <mike.rudenko@gmail.com>
Date:   Sun Oct 10 19:54:57 2021 +0200

    media: rockchip: rkisp1: use device name for debugfs subdir name

    [ Upstream commit c2611e479f5d9b05108270e1ab423955486b4457 ]

    While testing Rockchip RK3399 with both ISPs enabled, a dmesg error
    was observed:
    ```
    [15.559141] debugfs: Directory 'rkisp1' with parent '/' already present!
    ```

    Fix it by using the device name for the debugfs subdirectory name
    instead of the driver name, thus preventing name collision.

    Link: https://lore.kernel.org/linux-media/20211010175457.438627-1-mike.rudenko@gmail.com
    Signed-off-by: Mikhail Rudenko <mike.rudenko@gmail.com>
    Reviewed-by: Ezequiel Garcia <ezequiel@vanguardiasur.com.ar>
    Reviewed-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a2eec0402e7d6feeda580e0dc4e1b12779fdce7d
Author: Sean Young <sean@mess.org>
Date:   Tue Nov 30 23:58:19 2021 +0100

    media: igorplugusb: receiver overflow should be reported

    [ Upstream commit 8fede658e7ddb605bbd68ed38067ddb0af033db4 ]

    Without this, some IR will be missing mid-stream and we might decode
    something which never really occurred.
```

```
    Signed-off-by: Sean Young <sean@mess.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ddf6110d0592ab8274d2672d581cfc52513bfc10
Author: Alistair Francis <alistair@alistair23.me>
Date:   Wed Dec 8 22:40:44 2021 +1000

    HID: i2c-hid-of: Expose the touchscreen-inverted properties

    [ Upstream commit b60d3c803d7603432a08aeaf988aff53b3a5ec64 ]

    Allow the touchscreen-inverted-x/y device tree properties to control the
    HID_QUIRK_X_INVERT/HID_QUIRK_Y_INVERT quirks for the hid-input device.

    Signed-off-by: Alistair Francis <alistair@alistair23.me>
    Acked-by: Rob Herring <robh@kernel.org>
    [bentiss: silence checkpatch warnings]
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Link: https://lore.kernel.org/r/20211208124045.61815-3-alistair@alistair23.me
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8ef5b0350772a9c94506e5d0c2390daad1c9a5a
Author: Alistair Francis <alistair@alistair23.me>
Date:   Wed Dec 8 22:40:43 2021 +1000

    HID: quirks: Allow inverting the absolute X/Y values

    [ Upstream commit fd8d135b2c5e88662f2729e034913f183455a667 ]

    Add a HID_QUIRK_X_INVERT/HID_QUIRK_Y_INVERT quirk that can be used
    to invert the X/Y values.

    Signed-off-by: Alistair Francis <alistair@alistair23.me>
    [bentiss: silence checkpatch warning]
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Link: https://lore.kernel.org/r/20211208124045.61815-2-alistair@alistair23.me
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 781b39c3c6b445a15581037c9be00b767b2f7fa6
Author: Isabella Basso <isabbasso@riseup.net>
Date:   Thu Dec 9 12:47:19 2021 -0300

    drm/amdgpu: fix amdgpu_ras_mca_query_error_status scope

    [ Upstream commit 929bb8e200412da36aca4b61209ec26283f9c184 ]

    This commit fixes the compile-time warning below:

     warning: no previous prototype for ‘amdgpu_ras_mca_query_error_status’
     [-Wmissing-prototypes]

    Changes since v1:
    - As suggested by Alexander Deucher:
      1. Make function static instead of adding prototype.

    Signed-off-by: Isabella Basso <isabbasso@riseup.net>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 174a76f3e522c192c602a54731af5cb18b994a3e
Author: Felix Kuehling <Felix.Kuehling@amd.com>
Date:   Tue Dec 7 20:23:14 2021 -0500

    drm/amdkfd: Fix error handling in svm_range_add
```

```
    [ Upstream commit 726be40607264b180a2b336c81e1dcff941de618 ]

    Add null-pointer check after the last svm_range_new call. This was
    originally reported by Zhou Qingyang <zhou1615@umn.edu> based on a
    static analyzer.

    To avoid duplicating the unwinding code from svm_range_handle_overlap,
    I merged the two functions into one.

    Signed-off-by: Felix Kuehling <Felix.Kuehling@amd.com>
    Cc: Zhou Qingyang <zhou1615@umn.edu>
    Reviewed-by: Philip Yang <Philip.Yang@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1a4cefba5cb2fc9d74e859b2e8b21fac97e5563
Author: Paolo Abeni <pabeni@redhat.com>
Date:   Tue Nov 30 11:08:06 2021 +0100

    bpf: Do not WARN in bpf_warn_invalid_xdp_action()

    [ Upstream commit 2cbad989033bff0256675c38f96f5faab852af4b ]

    The WARN_ONCE() in bpf_warn_invalid_xdp_action() can be triggered by
    any bugged program, and even attaching a correct program to a NIC
    not supporting the given action.

    The resulting splat, beyond polluting the logs, fouls automated tools:
    e.g. a syzkaller reproducers using an XDP program returning an
    unsupported action will never pass validation.

    Replace the WARN_ONCE with a less intrusive pr_warn_once().

    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Toke Høiland-Jørgensen <toke@redhat.com>
    Link:
https://lore.kernel.org/bpf/016ceec56e4817ebb2a9e35ce794d5c917df572c.1638189075.git.pabeni@r
edhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4e23af09ffc9b88e1857c084709be674ecf3f7e
Author: David Gow <davidgow@google.com>
Date:   Tue Nov 2 00:30:13 2021 -0700

    kunit: Don't crash if no parameters are generated

    [ Upstream commit 37dbb4c7c7442dbfc9b651e4ddd4afe30b26afc9 ]

    It's possible that a parameterised test could end up with zero
    parameters. At the moment, the test function will nevertheless be called
    with NULL as the parameter. Instead, don't try to run the test code, and
    just mark the test as SKIPped.

    Reported-by: Daniel Latypov <dlatypov@google.com>
    Signed-off-by: David Gow <davidgow@google.com>
    Reviewed-by: Daniel Latypov <dlatypov@google.com>
    Reviewed-by: Brendan Higgins <brendanhiggins@google.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d3d46042a27634933c2c2e6b411f7af3e782ba33
Author: Suresh Kumar <surkumar@redhat.com>
Date:   Mon Dec 13 11:17:09 2021 +0530
```

```
    net: bonding: debug: avoid printing debug logs when bond is not notifying peers

    [ Upstream commit fee32de284ac277ba434a2d59f8ce46528ff3946 ]

    Currently "bond_should_notify_peers: slave ..." messages are printed whenever
    "bond_should_notify_peers" function is called.

    +++
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:26 node1 kernel: bond0: (slave enp0s25): Received LACPDU on port 1
    Dec 12 12:33:26 node1 kernel: bond0: (slave enp0s25): Rx Machine: Port=1, Last State=6,
Curr State=6
    Dec 12 12:33:26 node1 kernel: bond0: (slave enp0s25): partner sync=1
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:26 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    ...
    Dec 12 12:33:30 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:30 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:30 node1 kernel: bond0: (slave enp4s3): Received LACPDU on port 2
    Dec 12 12:33:30 node1 kernel: bond0: (slave enp4s3): Rx Machine: Port=2, Last State=6,
Curr State=6
    Dec 12 12:33:30 node1 kernel: bond0: (slave enp4s3): partner sync=1
    Dec 12 12:33:30 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:30 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    Dec 12 12:33:30 node1 kernel: bond0: bond_should_notify_peers: slave enp0s25
    +++

    This is confusing and can also clutter up debug logs.
    Print logs only when the peer notification happens.

    Signed-off-by: Suresh Kumar <suresh2514@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 657c53b65f9fea1c2c5e2a1cfe0dc706be65162d
Author: Borislav Petkov <bp@suse.de>
Date:   Tue Nov 2 11:14:48 2021 +0100

    x86/mce: Mark mce_read_aux() noinstr

    [ Upstream commit db6c996d6ce45dfb44891f0824a65ecec216f47a ]

    Fixes

      vmlinux.o: warning: objtool: do_machine_check()+0x681: call to mce_read_aux() leaves
.noinstr.text section

    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211208111343.8130-10-bp@alien8.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f11e312b0dc8ff5d9e135bb6b8a3007bab81ad4b
Author: Borislav Petkov <bp@suse.de>
Date:   Mon Nov 1 16:43:33 2021 +0100

    x86/mce: Mark mce_end() noinstr

    [ Upstream commit b4813539d37fa31fed62cdfab7bd2dd8929c5b2e ]

    It is called by the #MC handler which is noinstr.
```

```
    Fixes

      vmlinux.o: warning: objtool: do_machine_check()+0xbd6: call to memset() leaves
.noinstr.text section

    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211208111343.8130-9-bp@alien8.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf621be2fcabdd3d347995159907588b3d6be953
Author: Borislav Petkov <bp@suse.de>
Date:   Mon Nov 1 13:39:35 2021 +0100

    x86/mce: Mark mce_panic() noinstr

    [ Upstream commit 3c7ce80a818fa7950be123cac80cd078e5ac1013 ]

    And allow instrumentation inside it because it does calls to other
    facilities which will not be tagged noinstr.

    Fixes

      vmlinux.o: warning: objtool: do_machine_check()+0xc73: call to mce_panic() leaves
.noinstr.text section

    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211208111343.8130-8-bp@alien8.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9d8885178ce1444e799e98211da909515326648
Author: Borislav Petkov <bp@suse.de>
Date:   Wed Oct 13 09:47:50 2021 +0200

    x86/mce: Prevent severity computation from being instrumented

    [ Upstream commit 0a5b288e85bbef5227bb6397e31fcf1d7ba9142a ]

    Mark all the MCE severity computation logic noinstr and allow
    instrumentation when it "calls out".

    Fixes

      vmlinux.o: warning: objtool: do_machine_check()+0xc5d: call to mce_severity() leaves
.noinstr.text section

    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211208111343.8130-7-bp@alien8.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c55b322c2c98770be50c88bf52c5aaee058c3d1
Author: Borislav Petkov <bp@suse.de>
Date:   Wed Oct 13 09:07:19 2021 +0200

    x86/mce: Allow instrumentation during task work queueing

    [ Upstream commit 4fbce464db81a42f9a57ee242d6150ec7f996415 ]

    Fixes

      vmlinux.o: warning: objtool: do_machine_check()+0xdb1: call to queue_task_work()
leaves .noinstr.text section

    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211208111343.8130-6-bp@alien8.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit da582537133cac06536c26fef3c9e7d5eb79a363
Author: Alex Elder <elder@linaro.org>
Date:   Fri Dec 10 16:31:22 2021 -0600

    ARM: dts: qcom: sdx55: fix IPA interconnect definitions

    [ Upstream commit c0d6316c238b1bd743108bd4b08eda364f47c7c9 ]

    The first two interconnects defined for IPA on the SDX55 SoC are
    really two parts of what should be represented as a single path
    between IPA and system memory.

    Fix this by combining the "memory-a" and "memory-b" interconnects
    into a single "memory" interconnect.

    Reported-by: David Heidelberg <david@ixit.cz>
    Tested-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Alex Elder <elder@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 645a01aba88ad4beff1ad0d20be0640f747554b6
Author: Baochen Qiang <quic_bqiang@quicinc.com>
Date:   Thu Dec 9 09:19:49 2021 +0800

    ath11k: Avoid false DEADLOCK warning reported by lockdep

    [ Upstream commit 767c94caf0efad136157110787fe221b74cb5c8a ]

    With CONFIG_LOCKDEP=y and CONFIG_DEBUG_SPINLOCK=y, lockdep reports
    below warning:

    [  166.059415] ============================================
    [  166.059416] WARNING: possible recursive locking detected
    [  166.059418] 5.15.0-wt-ath+ #10 Tainted: G        W  O
    [  166.059420] --------------------------------------------
    [  166.059421] kworker/0:2/116 is trying to acquire lock:
    [  166.059423] ffff9905f2083160 (&srng->lock){+.-.}-{2:2}, at:
ath11k_hal_reo_cmd_send+0x20/0x490 [ath11k]
    [  166.059440]
                   but task is already holding lock:
    [  166.059442] ffff9905f2083230 (&srng->lock){+.-.}-{2:2}, at:
ath11k_dp_process_reo_status+0x95/0x2d0 [ath11k]
    [  166.059491]
                   other info that might help us debug this:
    [  166.059492]  Possible unsafe locking scenario:

    [  166.059493]        CPU0
    [  166.059494]        ----
    [  166.059495]   lock(&srng->lock);
    [  166.059498]   lock(&srng->lock);
    [  166.059500]
                    *** DEADLOCK ***

    [  166.059501]  May be due to missing lock nesting notation

    [  166.059502] 3 locks held by kworker/0:2/116:
    [  166.059504]  #0: ffff9905c0081548 ((wq_completion)events){+.+.}-{0:0}, at:
process_one_work+0x1f6/0x660
    [  166.059511]  #1: ffff9d2400a5fe68 ((debug_obj_work).work){+.+.}-{0:0}, at:
process_one_work+0x1f6/0x660
    [  166.059517]  #2: ffff9905f2083230 (&srng->lock){+.-.}-{2:2}, at:
ath11k_dp_process_reo_status+0x95/0x2d0 [ath11k]
    [  166.059532]
```

```
                  stack backtrace:
    [  166.059534] CPU: 0 PID: 116 Comm: kworker/0:2 Kdump: loaded Tainted: G        W  O
 5.15.0-wt-ath+ #10
    [  166.059537] Hardware name: Intel(R) Client Systems NUC8i7HVK/NUC8i7HVB, BIOS
 HNKBLi70.86A.0059.2019.1112.1124 11/12/2019
    [  166.059539] Workqueue: events free_obj_work
    [  166.059543] Call Trace:
    [  166.059545]  <IRQ>
    [  166.059547]  dump_stack_lvl+0x56/0x7b
    [  166.059552]  __lock_acquire+0xb9a/0x1a50
    [  166.059556]  lock_acquire+0x1e2/0x330
    [  166.059560]  ? ath11k_hal_reo_cmd_send+0x20/0x490 [ath11k]
    [  166.059571]  _raw_spin_lock_bh+0x33/0x70
    [  166.059574]  ? ath11k_hal_reo_cmd_send+0x20/0x490 [ath11k]
    [  166.059584]  ath11k_hal_reo_cmd_send+0x20/0x490 [ath11k]
    [  166.059594]  ath11k_dp_tx_send_reo_cmd+0x3f/0x130 [ath11k]
    [  166.059605]  ath11k_dp_rx_tid_del_func+0x221/0x370 [ath11k]
    [  166.059618]  ath11k_dp_process_reo_status+0x22f/0x2d0 [ath11k]
    [  166.059632]  ? ath11k_dp_service_srng+0x2ea/0x2f0 [ath11k]
    [  166.059643]  ath11k_dp_service_srng+0x2ea/0x2f0 [ath11k]
    [  166.059655]  ath11k_pci_ext_grp_napi_poll+0x1c/0x70 [ath11k_pci]
    [  166.059659]  __napi_poll+0x28/0x230
    [  166.059664]  net_rx_action+0x285/0x310
    [  166.059668]  __do_softirq+0xe6/0x4d2
    [  166.059672]  irq_exit_rcu+0xd2/0xf0
    [  166.059675]  common_interrupt+0xa5/0xc0
    [  166.059678]  </IRQ>
    [  166.059679]  <TASK>
    [  166.059680]  asm_common_interrupt+0x1e/0x40
    [  166.059683] RIP: 0010:_raw_spin_unlock_irqrestore+0x38/0x70
    [  166.059686] Code: 83 c7 18 e8 2a 95 43 ff 48 89 ef e8 22 d2 43 ff 81 e3 00 02 00 00
 75 25 9c 58 f6 c4 02 75 2d 48 85 db 74 01 fb bf 01 00 00 00 <e8> 63 2e 40 ff 65 8b 05 8c 59
 97 5c 85 c0 74 0a 5b 5d c3 e8 00 6a
    [  166.059689] RSP: 0018:ffff9d2400a5fca0 EFLAGS: 00000206
    [  166.059692] RAX: 0000000000000002 RBX: 0000000000000200 RCX: 0000000000000006
    [  166.059694] RDX: 0000000000000000 RSI: ffffffffa404879b RDI: 0000000000000001
    [  166.059696] RBP: ffff9905c0053000 R08: 0000000000000001 R09: 0000000000000001
    [  166.059698] R10: ffff9d2400a5fc50 R11: 0000000000000001 R12: ffffe186c41e2840
    [  166.059700] R13: 0000000000000001 R14: ffff9905c78a1c68 R15: 0000000000000001
    [  166.059704]  free_debug_processing+0x257/0x3d0
    [  166.059708]  ? free_obj_work+0x1f5/0x250
    [  166.059712]  __slab_free+0x374/0x5a0
    [  166.059718]  ? kmem_cache_free+0x2e1/0x370
    [  166.059721]  ? free_obj_work+0x1f5/0x250
    [  166.059724]  kmem_cache_free+0x2e1/0x370
    [  166.059727]  free_obj_work+0x1f5/0x250
    [  166.059731]  process_one_work+0x28b/0x660
    [  166.059735]  ? process_one_work+0x660/0x660
    [  166.059738]  worker_thread+0x37/0x390
    [  166.059741]  ? process_one_work+0x660/0x660
    [  166.059743]  kthread+0x176/0x1a0
    [  166.059746]  ? set_kthread_struct+0x40/0x40
    [  166.059749]  ret_from_fork+0x22/0x30
    [  166.059754]  </TASK>

    Since these two lockes are both initialized in ath11k_hal_srng_setup,
    they are assigned with the same key. As a result lockdep suspects that
    the task is trying to acquire the same lock (due to same key) while
    already holding it, and thus reports the DEADLOCK warning. However as
    they are different spinlock instances, the warning is false positive.

    On the other hand, even no dead lock indeed, this is a major issue for
    upstream regression testing as it disables lockdep functionality.

    Fix it by assigning separate lock class key for each srng->lock.
```

```
    Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1
    Signed-off-by: Baochen Qiang <quic_bqiang@quicinc.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20211209011949.151472-1-quic_bqiang@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3d47177d1897e1d914002f5161a57826e723f0b6
Author: Jagan Teki <jagan@amarulasolutions.com>
Date:   Fri Nov 12 19:53:59 2021 +0530

    arm64: dts: rockchip: Fix Bluetooth on ROCK Pi 4 boards

    [ Upstream commit f471b1b2db0819917c54099ab68349ad6a7e9e19 ]

    This patch fixes the Bluetooth on ROCK Pi 4 boards.

    ROCK Pi 4 boards has BCM4345C5 and now it is supported
    on Mainline Linux, brcm,bcm43438-bt still working but
    observed the BT Audio issues with latest test.

    So, use the BCM4345C5 compatible and its associated
    properties like clock-names as lpo and max-speed.

    Attach vbat and vddio supply rails as well.

    Signed-off-by: Jagan Teki <jagan@amarulasolutions.com>
    Link: https://lore.kernel.org/r/20211112142359.320798-1-jagan@amarulasolutions.com
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f22d581e055174e887709c473b0bc64d91f0d4eb
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Tue Nov 30 13:25:46 2021 +0100

    selftests/ftrace: make kprobe profile testcase description unique

    [ Upstream commit e5992f373c6eed6d09e5858e9623df1259b3ce30 ]

    Commit 32f6e5da83c7 ("selftests/ftrace: Add kprobe profile testcase")
    added a new kprobes testcase, but has a description which does not
    describe what the test case is doing and is duplicating the description
    of another test case.

    Therefore change the test case description, so it is unique and then
    allows easily to tell which test case actually passed or failed.

    Reported-by: Alexander Egorenkov <egorenar@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1170233f063de7eb9020344338966fd30afea7b6
Author: Iwona Winiarska <iwona.winiarska@intel.com>
Date:   Sat Dec 4 18:10:27 2021 +0100

    gpio: aspeed-sgpio: Convert aspeed_sgpio.lock to raw_spinlock

    [ Upstream commit ab39d6988dd53f354130438d8afa5596a2440fed ]

    The gpio-aspeed-sgpio driver implements an irq_chip which need to be
    invoked from hardirq context. Since spin_lock() can sleep with
    PREEMPT_RT, it is no longer legal to invoke it while interrupts are
    disabled.
```

```
    This also causes lockdep to complain about:
    [   25.919465] [ BUG: Invalid wait context ]
    because aspeed_sgpio.lock (spin_lock_t) is taken under irq_desc.lock
    (raw_spinlock_t).
    Let's use of raw_spinlock_t instead of spinlock_t.

    Signed-off-by: Iwona Winiarska <iwona.winiarska@intel.com>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 102dd8b25de70d91db4736b3e11b9a44308babd2
Author: Iwona Winiarska <iwona.winiarska@intel.com>
Date:   Sat Dec 4 18:10:26 2021 +0100

    gpio: aspeed: Convert aspeed_gpio.lock to raw_spinlock

    [ Upstream commit 61a7904b6ace99b1bde0d0e867fa3097f5c8cee2 ]

    The gpio-aspeed driver implements an irq_chip which need to be invoked
    from hardirq context. Since spin_lock() can sleep with PREEMPT_RT, it is
    no longer legal to invoke it while interrupts are disabled.
    This also causes lockdep to complain about:
    [    0.649797] [ BUG: Invalid wait context ]
    because aspeed_gpio.lock (spin_lock_t) is taken under irq_desc.lock
    (raw_spinlock_t).
    Let's use of raw_spinlock_t instead of spinlock_t.

    Signed-off-by: Iwona Winiarska <iwona.winiarska@intel.com>
    Signed-off-by: Bartosz Golaszewski <brgl@bgdev.pl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 23f32f364bfd5ed042054fb7a4e5de4bae28745b
Author: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
Date:   Wed Dec 8 11:36:30 2021 +0000

    net: phy: prefer 1000baseT over 1000baseKX

    [ Upstream commit f20f94f7f52c4685c81754f489ffcc72186e8bdb ]

    The PHY settings table is supposed to be sorted by descending match
    priority - in other words, earlier entries are preferred over later
    entries.

    The order of 1000baseKX/Full and 1000baseT/Full is such that we
    prefer 1000baseKX/Full over 1000baseT/Full, but 1000baseKX/Full is
    a lot rarer than 1000baseT/Full, and thus is much less likely to
    be preferred.

    This causes phylink problems - it means a fixed link specifying a
    speed of 1G and full duplex gets an ethtool linkmode of 1000baseKX/Full
    rather than 1000baseT/Full as would be expected - and since we offer
    userspace a software emulation of a conventional copper PHY, we want
    to offer copper modes in preference to anything else. However, we do
    still want to allow the rarer modes as well.

    Hence, let's reorder these two modes to prefer copper.

    Tested-by: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Reported-by: Florian Fainelli <f.fainelli@gmail.com>
    Link: https://lore.kernel.org/r/E1muvFO-00F6jY-1K@rmk-PC.armlinux.org.uk
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit af09ea6aab04d843858f744c55558f85c10b2bc7
Author: Wen Gong <quic_wgong@quicinc.com>
Date:   Wed Dec 8 01:17:52 2021 -0500

    ath10k: drop beacon and probe response which leak from other channel

    [ Upstream commit 3bf2537ec2e33310b431b53fd84be8833736c256 ]

    When scan request on channel 1, it also receive beacon from other
    channels, and the beacon also indicate to mac80211 and wpa_supplicant,
    and then the bss info appears in radio measurement report of radio
    measurement sent from wpa_supplicant, thus lead RRM case fail.

    This is to drop the beacon and probe response which is not the same
    channel of scanning.

    Tested-on: QCA6174 hw3.2 SDIO WLAN.RMH.4.4.1-00049

    Signed-off-by: Wen Gong <quic_wgong@quicinc.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20211208061752.16564-1-quic_wgong@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c5eb468cbc1fa663bf0cc6c5360802dea4e611c2
Author: Antoine Tenart <atenart@kernel.org>
Date:   Tue Dec 7 15:57:24 2021 +0100

    net-sysfs: update the queue counts in the unregistration path

    [ Upstream commit d7dac083414eb5bb99a6d2ed53dc2c1b405224e5 ]

    When updating Rx and Tx queue kobjects, the queue count should always be
    updated to match the queue kobjects count. This was not done in the net
    device unregistration path, fix it. Tracking all queue count updates
    will allow in a following up patch to detect illegal updates.

    Signed-off-by: Antoine Tenart <atenart@kernel.org>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce7eff571f05a3ff4de91d84a79feb3a13c16c10
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Wed Dec 1 17:38:16 2021 +0800

    rtw89: don't kick off TX DMA if failed to write skb

    [ Upstream commit a58fdb7c843a37d6598204c6513961feefdadc6a ]

    This is found by Smatch static checker warning:
            drivers/net/wireless/realtek/rtw89/mac80211.c:31 rtw89_ops_tx()
            error: uninitialized symbol 'qsel'.

    The warning is because 'qsel' isn't filled by rtw89_core_tx_write() due to
    failed to write. The way to fix it is to avoid kicking off TX DMA, so add
    'return' to the failure case.

    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20211201093816.13806-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d5d7ff21515035e0eb1c485eab8be1cdac8fbb2
Author: Sebastian Gottschall <s.gottschall@dd-wrt.com>
Date:   Wed May 5 15:58:06 2021 +0700
```

```
    ath10k: Fix tx hanging

    [ Upstream commit e8a91863eba3966a447d2daa1526082d52b5db2a ]

    While running stress tests in roaming scenarios (switching ap's every 5
    seconds, we discovered a issue which leads to tx hangings of exactly 5
    seconds while or after scanning for new accesspoints. We found out that
    this hanging is triggered by ath10k_mac_wait_tx_complete since the
    empty_tx_wq was not wake when the num_tx_pending counter reaches zero.
    To fix this, we simply move the wake_up call to htt_tx_dec_pending,
    since this call was missed on several locations within the ath10k code.

    Signed-off-by: Sebastian Gottschall <s.gottschall@dd-wrt.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20210505085806.11474-1-s.gottschall@dd-wrt.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3fc384692ea634855e047072e5d53d78cd2bd112
Author: Wen Gong <quic_wgong@quicinc.com>
Date:   Tue Dec 7 17:23:36 2021 +0200

    ath11k: avoid deadlock by change ieee80211_queue_work for regd_update_work

    [ Upstream commit ed05c7cf1286d7e31e7623bce55ff135723591bf ]

    When enable debug config, it print below warning while shut down wlan
    interface shuh as run "ifconfig wlan0 down".

    The reason is because ar->regd_update_work is ran once, and it is will
    call wiphy_lock(ar->hw->wiphy) in function ath11k_regd_update() which
    is running in workqueue of ieee80211_local queued by ieee80211_queue_work().
    Another thread from "ifconfig wlan0 down" will also accuqire the lock
    by wiphy_lock(sdata->local->hw.wiphy) in function ieee80211_stop(), and
    then it call ieee80211_stop_device() to flush_workqueue(local->workqueue),
    this will wait the workqueue of ieee80211_local finished. Then deadlock
    will happen easily if the two thread run meanwhile.

    Below warning disappeared after this change.

    [  914.088798] ath11k_pci 0000:05:00.0: mac remove interface (vdev 0)
    [  914.088806] ath11k_pci 0000:05:00.0: mac stop 11d scan
    [  914.088810] ath11k_pci 0000:05:00.0: mac stop 11d vdev id 0
    [  914.088827] ath11k_pci 0000:05:00.0: htc ep 2 consumed 1 credits (total 0)
    [  914.088841] ath11k_pci 0000:05:00.0: send 11d scan stop vdev id 0
    [  914.088849] ath11k_pci 0000:05:00.0: htc insufficient credits ep 2 required 1
available 0
    [  914.088856] ath11k_pci 0000:05:00.0: htc insufficient credits ep 2 required 1
available 0
    [  914.096434] ath11k_pci 0000:05:00.0: rx ce pipe 2 len 16
    [  914.096442] ath11k_pci 0000:05:00.0: htc ep 2 got 1 credits (total 1)
    [  914.096481] ath11k_pci 0000:05:00.0: htc ep 2 consumed 1 credits (total 0)
    [  914.096491] ath11k_pci 0000:05:00.0: WMI vdev delete id 0
    [  914.111598] ath11k_pci 0000:05:00.0: rx ce pipe 2 len 16
    [  914.111628] ath11k_pci 0000:05:00.0: htc ep 2 got 1 credits (total 1)
    [  914.114659] ath11k_pci 0000:05:00.0: rx ce pipe 2 len 20
    [  914.114742] ath11k_pci 0000:05:00.0: htc rx completion ep 2 skb         pK-error
    [  914.115977] ath11k_pci 0000:05:00.0: vdev delete resp for vdev id 0
    [  914.116685] ath11k_pci 0000:05:00.0: vdev 00:03:7f:29:61:11 deleted, vdev_id 0

    [  914.117583] =====================================================
    [  914.117592] WARNING: possible circular locking dependency detected
    [  914.117600] 5.16.0-rc1-wt-ath+ #1 Tainted: G           OE
    [  914.117611] -----------------------------------------------------
    [  914.117618] ifconfig/2805 is trying to acquire lock:
```

```
     [  914.117628] ffff9c00a62bb548 ((wq_completion)phy0){+.+.}-{0:0}, at:
  flush_workqueue+0x87/0x470
     [  914.117674]
                    but task is already holding lock:
     [  914.117682] ffff9c00baea07d0 (&rdev->wiphy.mtx){+.+.}-{4:4}, at:
  ieee80211_stop+0x38/0x180 [mac80211]
     [  914.117872]
                    which lock already depends on the new lock.

     [  914.117880]
                    the existing dependency chain (in reverse order) is:
     [  914.117888]
                    -> #3 (&rdev->wiphy.mtx){+.+.}-{4:4}:
     [  914.117910]         __mutex_lock+0xa0/0x9c0
     [  914.117930]         mutex_lock_nested+0x1b/0x20
     [  914.117944]         reg_process_self_managed_hints+0x3a/0xb0 [cfg80211]
     [  914.118093]         wiphy_regulatory_register+0x47/0x80 [cfg80211]
     [  914.118229]         wiphy_register+0x84f/0x9c0 [cfg80211]
     [  914.118353]         ieee80211_register_hw+0x6b1/0xd90 [mac80211]
     [  914.118486]         ath11k_mac_register+0x6af/0xb60 [ath11k]
     [  914.118550]         ath11k_core_qmi_firmware_ready+0x383/0x4a0 [ath11k]
     [  914.118598]         ath11k_qmi_driver_event_work+0x347/0x4a0 [ath11k]
     [  914.118656]         process_one_work+0x228/0x670
     [  914.118669]         worker_thread+0x4d/0x440
     [  914.118680]         kthread+0x16d/0x1b0
     [  914.118697]         ret_from_fork+0x22/0x30
     [  914.118714]
                    -> #2 (rtnl_mutex){+.+.}-{4:4}:
     [  914.118736]         __mutex_lock+0xa0/0x9c0
     [  914.118751]         mutex_lock_nested+0x1b/0x20
     [  914.118767]         rtnl_lock+0x17/0x20
     [  914.118783]         ath11k_regd_update+0x15a/0x260 [ath11k]
     [  914.118841]         ath11k_regd_update_work+0x15/0x20 [ath11k]
     [  914.118897]         process_one_work+0x228/0x670
     [  914.118909]         worker_thread+0x4d/0x440
     [  914.118920]         kthread+0x16d/0x1b0
     [  914.118934]         ret_from_fork+0x22/0x30
     [  914.118948]
                    -> #1 ((work_completion)(&ar->regd_update_work)){+.+.}-{0:0}:
     [  914.118972]         process_one_work+0x1fa/0x670
     [  914.118984]         worker_thread+0x4d/0x440
     [  914.118996]         kthread+0x16d/0x1b0
     [  914.119010]         ret_from_fork+0x22/0x30
     [  914.119023]
                    -> #0 ((wq_completion)phy0){+.+.}-{0:0}:
     [  914.119045]         __lock_acquire+0x146d/0x1cf0
     [  914.119057]         lock_acquire+0x19b/0x360
     [  914.119067]         flush_workqueue+0xae/0x470
     [  914.119084]         ieee80211_stop_device+0x3b/0x50 [mac80211]
     [  914.119260]         ieee80211_do_stop+0x5d7/0x830 [mac80211]
     [  914.119409]         ieee80211_stop+0x45/0x180 [mac80211]
     [  914.119557]         __dev_close_many+0xb3/0x120
     [  914.119573]         __dev_change_flags+0xc3/0x1d0
     [  914.119590]         dev_change_flags+0x29/0x70
     [  914.119605]         devinet_ioctl+0x653/0x810
     [  914.119620]         inet_ioctl+0x193/0x1e0
     [  914.119631]         sock_do_ioctl+0x4d/0xf0
     [  914.119649]         sock_ioctl+0x262/0x340
     [  914.119665]         __x64_sys_ioctl+0x96/0xd0
     [  914.119678]         do_syscall_64+0x3d/0xd0
     [  914.119694]         entry_SYSCALL_64_after_hwframe+0x44/0xae
     [  914.119709]
                    other info that might help us debug this:

     [  914.119717] Chain exists of:
```

```
                    (wq_completion)phy0 --> rtnl_mutex --> &rdev->wiphy.mtx

    [  914.119745]  Possible unsafe locking scenario:

    [  914.119752]        CPU0                    CPU1
    [  914.119758]        ----                    ----
    [  914.119765]   lock(&rdev->wiphy.mtx);
    [  914.119778]                                lock(rtnl_mutex);
    [  914.119792]                                lock(&rdev->wiphy.mtx);
    [  914.119807]   lock((wq_completion)phy0);
    [  914.119819]
                   *** DEADLOCK ***

    [  914.119827] 2 locks held by ifconfig/2805:
    [  914.119837]  #0: ffffffffba3dc010 (rtnl_mutex){+.+.}-{4:4}, at: rtnl_lock+0x17/0x20
    [  914.119872]  #1: ffff9c00baea07d0 (&rdev->wiphy.mtx){+.+.}-{4:4}, at:
ieee80211_stop+0x38/0x180 [mac80211]
    [  914.120039]
                  stack backtrace:
    [  914.120048] CPU: 0 PID: 2805 Comm: ifconfig Tainted: G           OE     5.16.0-rc1-
wt-ath+ #1
    [  914.120064] Hardware name: LENOVO 418065C/418065C, BIOS 83ET63WW (1.33 ) 07/29/2011
    [  914.120074] Call Trace:
    [  914.120084]  <TASK>
    [  914.120094]  dump_stack_lvl+0x73/0xa4
    [  914.120119]  dump_stack+0x10/0x12
    [  914.120135]  print_circular_bug.isra.44+0x221/0x2e0
    [  914.120165]  check_noncircular+0x106/0x150
    [  914.120203]  __lock_acquire+0x146d/0x1cf0
    [  914.120215]  ? __lock_acquire+0x146d/0x1cf0
    [  914.120245]  lock_acquire+0x19b/0x360
    [  914.120259]  ? flush_workqueue+0x87/0x470
    [  914.120286]  ? lockdep_init_map_type+0x6b/0x250
    [  914.120310]  flush_workqueue+0xae/0x470
    [  914.120327]  ? flush_workqueue+0x87/0x470
    [  914.120344]  ? lockdep_hardirqs_on+0xd7/0x150
    [  914.120391]  ieee80211_stop_device+0x3b/0x50 [mac80211]
    [  914.120565]  ? ieee80211_stop_device+0x3b/0x50 [mac80211]
    [  914.120736]  ieee80211_do_stop+0x5d7/0x830 [mac80211]
    [  914.120906]  ieee80211_stop+0x45/0x180 [mac80211]
    [  914.121060]  __dev_close_many+0xb3/0x120
    [  914.121081]  __dev_change_flags+0xc3/0x1d0
    [  914.121109]  dev_change_flags+0x29/0x70
    [  914.121131]  devinet_ioctl+0x653/0x810
    [  914.121149]  ? __might_fault+0x77/0x80
    [  914.121179]  inet_ioctl+0x193/0x1e0
    [  914.121194]  ? inet_ioctl+0x193/0x1e0
    [  914.121218]  ? __might_fault+0x77/0x80
    [  914.121238]  ? _copy_to_user+0x68/0x80
    [  914.121266]  sock_do_ioctl+0x4d/0xf0
    [  914.121283]  ? inet_stream_connect+0x60/0x60
    [  914.121297]  ? sock_do_ioctl+0x4d/0xf0
    [  914.121329]  sock_ioctl+0x262/0x340
    [  914.121347]  ? sock_ioctl+0x262/0x340
    [  914.121362]  ? exit_to_user_mode_prepare+0x13b/0x280
    [  914.121388]  ? syscall_enter_from_user_mode+0x20/0x50
    [  914.121416]  __x64_sys_ioctl+0x96/0xd0
    [  914.121430]  ? br_ioctl_call+0x90/0x90
    [  914.121445]  ? __x64_sys_ioctl+0x96/0xd0
    [  914.121465]  do_syscall_64+0x3d/0xd0
    [  914.121482]  entry_SYSCALL_64_after_hwframe+0x44/0xae
    [  914.121497] RIP: 0033:0x7f0ed051737b
    [  914.121513] Code: 0f 1e fa 48 8b 05 15 3b 0d 00 64 c7 00 26 00 00 00 48 c7 c0 ff ff
ff ff c3 66 0f 1f 44 00 00 f3 0f 1e fa b8 10 00 00 00 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48
8b 0d e5 3a 0d 00 f7 d8 64 89 01 48
```

```
     [  914.121527] RSP: 002b:00007fff7be38b98 EFLAGS: 00000202 ORIG_RAX: 0000000000000010
     [  914.121544] RAX: ffffffffffffffda RBX: 00007fff7be38ba0 RCX: 00007f0ed051737b
     [  914.121555] RDX: 00007fff7be38ba0 RSI: 0000000000008914 RDI: 0000000000000004
     [  914.121566] RBP: 00007fff7be38c60 R08: 000000000000000a R09: 0000000000000001
     [  914.121576] R10: 0000000000000000 R11: 0000000000000202 R12: 00000000fffffffe
     [  914.121586] R13: 0000000000000004 R14: 0000000000000000 R15: 0000000000000000
     [  914.121620]  </TASK>

     Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1

     Signed-off-by: Wen Gong <quic_wgong@quicinc.com>
     Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
     Link: https://lore.kernel.org/r/20211201071745.17746-2-quic_wgong@quicinc.com
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f9e072cb964c354941a7c0ab7234270c01f7c6fd
Author: Wander Lairson Costa <wander@redhat.com>
Date:   Wed Nov 10 11:37:45 2021 -0300

     rcutorture: Avoid soft lockup during cpu stall

     [ Upstream commit 5ff7c9f9d7e3e0f6db5b81945fa11b69d62f433a ]

     If we use the module stall_cpu option, we may get a soft lockup warning
     in case we also don't pass the stall_cpu_block option.

     Introduce the stall_no_softlockup option to avoid a soft lockup on
     cpu stall even if we don't use the stall_cpu_block option.

     Signed-off-by: Wander Lairson Costa <wander@redhat.com>
     Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ede66f560cffc03407d1c032fa53217d787e2531
Author: Miri Korenblit <miriam.rachel.korenblit@intel.com>
Date:   Sat Dec 4 13:10:50 2021 +0200

     iwlwifi: acpi: fix wgds rev 3 size

     [ Upstream commit dc276ffd0754e94080565c10b964f3c211879fdd ]

     The exact size of WGDS revision 3 was calculated using the wrong
     parameters. Fix it.

     Signed-off-by: Miri Korenblit <miriam.rachel.korenblit@intel.com>
     Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
     Link:
https://lore.kernel.org/r/iwlwifi.20211204130722.12c5b0cffe52.I7f342502f628f43a7e000189a6994
84bcef0f562@changeid
     Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 041eaa7356279a71ef7731e8157b91e0394b23ea
Author: Shaul Triebitz <shaul.triebitz@intel.com>
Date:   Sat Dec 4 13:10:47 2021 +0200

     iwlwifi: mvm: avoid clearing a just saved session protection id

     [ Upstream commit 8e967c137df3b236d2075f9538cb888129425d1a ]

     When scheduling a session protection the id is saved but
     then it may be cleared when calling iwl_mvm_te_clear_data
     (if a previous session protection is currently active).
     Fix it by saving the id after calling iwl_mvm_te_clear_data.
```

```
    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.20211204130722.b0743a588d14.I098fef6677d0dab3ef1b6183ed206
a10bab01eb2@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35a89cdd07a43c09013a832c720ae28c571613a1
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Sat Dec 4 08:35:45 2021 +0200

    iwlwifi: mvm: synchronize with FW after multicast commands

    [ Upstream commit db66abeea3aefed481391ecc564fb7b7fb31d742 ]

    If userspace installs a lot of multicast groups very quickly, then
    we may run out of command queue space as we send the updates in an
    asynchronous fashion (due to locking concerns), and the CPU can
    create them faster than the firmware can process them. This is true
    even when mac80211 has a work struct that gets scheduled.

    Fix this by synchronizing with the firmware after sending all those
    commands - outside of the iteration we can send a synchronous echo
    command that just has the effect of the CPU waiting for the prior
    asynchronous commands to finish. This also will cause fewer of the
    commands to be sent to the firmware overall, because the work will
    only run once when rescheduled multiple times while it's running.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=213649
    Suggested-by: Emmanuel Grumbach <emmanuel.grumbach@intel.com>
    Reported-by: Maximilian Ernestus <maximilian@ernestus.de>
    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.20211204083238.51aea5b79ea4.I88a44798efda16e9fe480fb3e9422
4931d311b29@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7fb36817e76be9a2c2156d8862f4abc21762dc7
Author: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
Date:   Thu Nov 4 22:40:25 2021 +0000

    arm64: dts: renesas: Fix thermal bindings

    [ Upstream commit 82ce79391d0ec25ec8aaae3c0617b71048ff0836 ]

    The binding node names for the thermal zones are not successfully
    validated by the dt-schemas.

    Fix the validation by changing from sensor-thermalN or thermal-sensor-N
    to sensorN-thermal.  Provide node labels of the form sensorN_thermal to
    ensure consistency with the other platform implementations.

    Signed-off-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Link: https://lore.kernel.org/r/20211104224033.3997504-1-
kieran.bingham+renesas@ideasonboard.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2eb2475322639d13390324dcf636b085c22f6a0
Author: Mika Westerberg <mika.westerberg@linux.intel.com>
Date:   Sun Nov 14 16:07:11 2021 +0200
```

```
    thunderbolt: Runtime PM activate both ends of the device link

    [ Upstream commit f3380cac0c0b3a6f49ab161e2a057c363962f48d ]

    If protocol tunnels are already up when the driver is loaded, for
    instance if the boot firmware implements connection manager of its own,
    runtime PM reference count of the consumer devices behind the tunnel
    might have been increased already before the device link is created but
    the supplier device runtime PM reference count is not. This leads to a
    situation where the supplier (the Thunderbolt driver) can runtime
    suspend even if it should not because the corresponding protocol tunnel
    needs to be up causing the devices to be removed from the corresponding
    native bus.

    Prevent this from happening by making both sides of the link runtime PM
    active briefly. The pm_runtime_put() for the consumer (PCIe
    root/downstream port, xHCI) then allows it to runtime suspend again but
    keeps the supplier runtime resumed the whole time it is runtime active.

    Signed-off-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Reviewed-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7e221dc8f23727e00a7fb6709b3318547a7c4d8
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Mon Dec 6 15:34:19 2021 +0100

    media: m920x: don't use stack on USB reads

    [ Upstream commit a2ab06d7c4d6bfd0b545a768247a70463e977e27 ]

    Using stack-allocated pointers for USB message data don't work.
    This driver is almost OK with that, except for the I2C read
    logic.

    Fix it by using a temporary read buffer, just like on all other
    calls to m920x_read().

    Link: https://lore.kernel.org/all/ccc99e48-de4f-045e-0fe4-61e3118e3f74@mida.se/
    Reported-by: rkardell@mida.se
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3af199c448765c45338112724de3a4d14c7183d
Author: Tsuchiya Yuto <kitakar@gmail.com>
Date:   Wed Dec 1 15:19:04 2021 +0100

    media: atomisp: fix "variable dereferenced before check 'asd'"

    [ Upstream commit ac56760a8bbb4e654b2fd54e5de79dd5d72f937d ]

    There are two occurrences where the variable 'asd' is dereferenced
    before check. Fix this issue by using the variable after the check.

    Link: https://lore.kernel.org/linux-media/20211122074122.GA6581@kili/

    Link: https://lore.kernel.org/linux-media/20211201141904.47231-1-kitakar@gmail.com
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a6486e63b279a5c346c023de04861a34abab18e
Author: Zhou Qingyang <zhou1615@umn.edu>
```

```
Date:   Tue Nov 30 17:25:49 2021 +0100

    media: saa7146: hexium_orion: Fix a NULL pointer dereference in hexium_attach()

    [ Upstream commit 348df8035301dd212e3cc2860efe4c86cb0d3303 ]

    In hexium_attach(dev, info), saa7146_vv_init() is called to allocate
    a new memory for dev->vv_data. In hexium_detach(), saa7146_vv_release()
    will be called and there is a dereference of dev->vv_data in
    saa7146_vv_release(), which could lead to a NULL pointer dereference
    on failure of saa7146_vv_init() according to the following logic.

    Both hexium_attach() and hexium_detach() are callback functions of
    the variable 'extension', so there exists a possible call chain directly
    from hexium_attach() to hexium_detach():

    hexium_attach(dev, info) -- fail to alloc memory to dev->vv_data
            |                               in saa7146_vv_init().
            |
            |
    hexium_detach() -- a dereference of dev->vv_data in saa7146_vv_release()

    Fix this bug by adding a check of saa7146_vv_init().

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_VIDEO_HEXIUM_ORION=m show no new warnings,
    and our static analyzer no longer warns about this code.

    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aea76a417f3ba94a51a8686b71eca0584334d29c
Author: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
Date:   Fri Nov 26 00:02:57 2021 +0100

    media: rcar-vin: Update format alignment constraints

    [ Upstream commit da6911f330d40cfe115a37249e47643eff555e82 ]

    This change fixes two issues with the size constraints for buffers.

    - There is no width alignment constraint for RGB formats. Prior to this
      change they were treated as YUV and as a result were more restricted
      than needed. Add a new check to differentiate between the two.

    - The minimum width and height supported is 5x2, not 2x4, this is an
      artifact from the driver's soc-camera days. Fix this incorrect
      assumption.

    Signed-off-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 0ed6bc5a2eeb5365f9d0c5a61deae81398b211f8
Author: James Hilliard <james.hilliard1@gmail.com>
Date:   Sun Nov 14 09:52:36 2021 +0100

    media: uvcvideo: Increase UVC_CTRL_CONTROL_TIMEOUT to 5 seconds.

    [ Upstream commit c8ed7d2f614cd8b315981d116c7a2fb01829500d ]

    Some uvc devices appear to require the maximum allowed USB timeout
    for GET_CUR/SET_CUR requests.

    So lets just bump the UVC control timeout to 5 seconds which is the
    same as the usb ctrl get/set defaults:
    USB_CTRL_GET_TIMEOUT 5000
    USB_CTRL_SET_TIMEOUT 5000

    It fixes the following runtime warnings:
       Failed to query (GET_CUR) UVC control 11 on unit 2: -110 (exp. 1).
       Failed to query (SET_CUR) UVC control 3 on unit 2: -110 (exp. 2).

    Signed-off-by: James Hilliard <james.hilliard1@gmail.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2954722f1daf957dbba989de2a5c54f761254678
Author: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
Date:   Mon Nov 29 03:10:46 2021 +0200

    drm: rcar-du: Fix CRTC timings when CMM is used

    [ Upstream commit f0ce591dc9a97067c6e783a2eaccd22c5476144d ]

    When the CMM is enabled, an offset of 25 pixels must be subtracted from
    the HDS (horizontal display start) and HDE (horizontal display end)
    registers. Fix the timings calculation, and take this into account in
    the mode validation.

    This fixes a visible horizontal offset in the image with VGA monitors.
    HDMI monitors seem to be generally more tolerant to incorrect timings,
    but may be affected too.

    Signed-off-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b4e9590a12c2cb2e1dea1979d942db3bc1e4aa3
Author: Joerg Roedel <jroedel@suse.de>
Date:   Thu Dec 2 16:32:25 2021 +0100

    x86/mm: Flush global TLB when switching to trampoline page-table

    [ Upstream commit 71d5049b053876afbde6c3273250b76935494ab2 ]

    Move the switching code into a function so that it can be re-used and
    add a global TLB flush. This makes sure that usage of memory which is
    not mapped in the trampoline page-table is reliably caught.

    Also move the clearing of CR4.PCIDE before the CR3 switch because the
    cr4_clear_bits() function will access data not mapped into the
    trampoline page-table.

    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211202153226.22946-4-joro@8bytes.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 75aed8bc8cedbe3fd6386779145fb9e18def41c6
Author: Xiongwei Song <sxwjean@gmail.com>
Date:   Tue Nov 16 21:10:33 2021 +0800

    floppy: Add max size check for user space request

    [ Upstream commit 545a32498c536ee152331cd2e7d2416aa0f20e01 ]

    We need to check the max request size that is from user space before
    allocating pages. If the request size exceeds the limit, return -EINVAL.
    This check can avoid the warning below from page allocator.

    WARNING: CPU: 3 PID: 16525 at mm/page_alloc.c:5344 current_gfp_context
include/linux/sched/mm.h:195 [inline]
    WARNING: CPU: 3 PID: 16525 at mm/page_alloc.c:5344 __alloc_pages+0x45d/0x500
mm/page_alloc.c:5356
    Modules linked in:
    CPU: 3 PID: 16525 Comm: syz-executor.3 Not tainted 5.15.0-syzkaller #0
    Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.14.0-2 04/01/2014
    RIP: 0010:__alloc_pages+0x45d/0x500 mm/page_alloc.c:5344
    Code: be c9 00 00 00 48 c7 c7 20 4a 97 89 c6 05 62 32 a7 0b 01 e8 74 9a 42 07 e9 6a ff
ff ff 0f 0b e9 a0 fd ff ff 40 80 e5 3f eb 88 <0f> 0b e9 18 ff ff ff 4c 89 ef 44 89 e6 45 31
ed e8 1e 76 ff ff e9
    RSP: 0018:ffffc90023b87850 EFLAGS: 00010246
    RAX: 0000000000000000 RBX: 1ffff92004770f0b RCX: dffffc0000000000
    RDX: 0000000000000000 RSI: 0000000000000033 RDI: 0000000000010cc1
    RBP: 0000000000000000 R08: 0000000000000000 R09: 0000000000000001
    R10: ffffffff81bb4686 R11: 0000000000000001 R12: ffffffff902c1960
    R13: 0000000000000033 R14: 0000000000000000 R15: ffff88804cf64a30
    FS:  0000000000000000(0000) GS:ffff88802cd00000(0063) knlGS:00000000f44b4b40
    CS:  0010 DS: 002b ES: 002b CR0: 0000000080050033
    CR2: 000000002c921000 CR3: 000000004f507000 CR4: 0000000000150ee0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     <TASK>
     alloc_pages+0x1a7/0x300 mm/mempolicy.c:2191
     __get_free_pages+0x8/0x40 mm/page_alloc.c:5418
     raw_cmd_copyin drivers/block/floppy.c:3113 [inline]
     raw_cmd_ioctl drivers/block/floppy.c:3160 [inline]
     fd_locked_ioctl+0x12e5/0x2820 drivers/block/floppy.c:3528
     fd_ioctl drivers/block/floppy.c:3555 [inline]
     fd_compat_ioctl+0x891/0x1b60 drivers/block/floppy.c:3869
     compat_blkdev_ioctl+0x3b8/0x810 block/ioctl.c:662
     __do_compat_sys_ioctl+0x1c7/0x290 fs/ioctl.c:972
     do_syscall_32_irqs_on arch/x86/entry/common.c:112 [inline]
     __do_fast_syscall_32+0x65/0xf0 arch/x86/entry/common.c:178
     do_fast_syscall_32+0x2f/0x70 arch/x86/entry/common.c:203
     entry_SYSENTER_compat_after_hwframe+0x4d/0x5c

    Reported-by: syzbot+23a02c7df2cf2bc93fa2@syzkaller.appspotmail.com
    Link: https://lore.kernel.org/r/20211116131033.27685-1-sxwjean@me.com
    Signed-off-by: Xiongwei Song <sxwjean@gmail.com>
    Signed-off-by: Denis Efremov <efremov@linux.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c5ca7a97eb4e89f38de88d904cdc6043ae078d85
Author: Neal Liu <neal_liu@aspeedtech.com>
Date:   Fri Nov 26 18:00:21 2021 +0800

    usb: uhci: add aspeed ast2600 uhci support

    [ Upstream commit 554abfe2eadec97d12c71d4a69da1518478f69eb ]
```

```
    Enable ast2600 uhci quirks.

    Signed-off-by: Neal Liu <neal_liu@aspeedtech.com>
    Link: https://lore.kernel.org/r/20211126100021.2331024-1-neal_liu@aspeedtech.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b128f148e344d1aeea4ef18debce9100f6f1b5b
Author: Kishon Vijay Abraham I <kishon@ti.com>
Date:   Fri Nov 26 14:15:55 2021 +0530

    arm64: dts: ti: j721e-main: Fix 'dtbs_check' in serdes_ln_ctrl node

    [ Upstream commit 3f92a5be6084b77f764a8bbb881ac0d12cb9e863 ]

    Fix 'dtbs_check' in serdes_ln_ctrl (mux@4080) node by changing the node
    name to mux-controller@4080.

    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211126084555.17797-3-kishon@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 567061fa04f2141ffb3010574ac658df805e6e2e
Author: Kishon Vijay Abraham I <kishon@ti.com>
Date:   Fri Nov 26 14:15:54 2021 +0530

    arm64: dts: ti: j7200-main: Fix 'dtbs_check' serdes_ln_ctrl node

    [ Upstream commit 4d3984906397581dc0ccb6a02bf16b6ff82c9192 ]

    Fix 'dtbs_check' in serdes_ln_ctrl (serdes-ln-ctrl@4080) node by
    changing the node name to mux-controller@4080.

    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211126084555.17797-2-kishon@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 746a39d26d00b22bbddc9ae3843649873655fda8
Author: Vlad Zahorodnii <vlad.zahorodnii@kde.org>
Date:   Thu Dec 2 14:52:15 2021 +0200

    drm/amd/display: Use oriented source size when checking cursor scaling

    [ Upstream commit 69cb56290d9d10cdcc461aa2685e67e540507a96 ]

    dm_check_crtc_cursor() doesn't take into account plane transforms when
    calculating plane scaling, this can result in false positives.

    For example, if there's an output with resolution 3840x2160 and the
    output is rotated 90 degrees, CRTC_W and CRTC_H will be 3840 and 2160,
    respectively, but SRC_W and SRC_H will be 2160 and 3840, respectively.

    Since the cursor plane usually has a square buffer attached to it, the
    dm_check_crtc_cursor() will think that there's a scale factor mismatch
    even though there isn't really.

    This fixes an issue where kwin fails to use hardware plane transforms.

    Changes since version 1:
    - s/orientated/oriented/g
```

```
    Reviewed-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Signed-off-by: Vlad Zahorodnii <vlad.zahorodnii@kde.org>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 657b5d916a60291d4590b349891503967b1be74b
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Tue Nov 30 10:52:55 2021 +0100

    drm: Return error codes from struct drm_driver.gem_create_object

    [ Upstream commit 4ff22f487f8c26b99cbe1678344595734c001a39 ]

    GEM helper libraries use struct drm_driver.gem_create_object to let
    drivers override GEM object allocation. On failure, the call returns
    NULL.

    Change the semantics to make the calls return a pointer-encoded error.
    This aligns the callback with its callers. Fixes the ingenic driver,
    which already returns an error pointer.

    Also update the callers to handle the involved types more strictly.

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Reviewed-by: Steven Price <steven.price@arm.com>
    Acked-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211130095255.26710-1-
tzimmermann@suse.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf7f10ace55440eb84356af4186353d52b9a008f
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon Nov 22 18:05:33 2021 +0100

    ACPI / x86: Add not-present quirk for the PCI0.SDHB.BRC1 device on the GPD win

    [ Upstream commit 57d2dbf710d832841872fb15ebb79429cab90fae ]

    The GPD win and its sibling the GPD pocket (99% the same electronics in a
    different case) use a PCI wifi card. But the ACPI tables on both variants
    contain a bug where the SDIO MMC controller for SDIO wifi cards is enabled
    despite this. This SDIO MMC controller has a PCI0.SDHB.BRC1 child-device
    which _PS3 method sets a GPIO causing the PCI wifi card to turn off.

    At the moment there is a pretty ugly kludge in the sdhci-acpi.c code,
    just to work around the bug in the DSDT of this single design. This can
    be solved cleaner/simply with a quirk overriding the _STA return of the
    broken PCI0.SDHB.BRC1 PCI0.SDHB.BRC1 child with a status value of 0,
    so that its power_manageable flag gets cleared, avoiding this problem.

    Note that even though it is not used, the _STA method for the MMC
    controller is deliberately not overridden. If the status of the MMC
    controller were forced to 0 it would never get suspended, which would
    cause these mini-laptops to not reach S0i3 level when suspended.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db54edf57cfb06e401c3012a96d526d6a292c5a7
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon Nov 22 18:05:32 2021 +0100

    ACPI / x86: Allow specifying acpi_device_override_status() quirks by path
```

```
    [ Upstream commit ba46e42e925b5d09b4e441f8de3db119cc7df58f ]

    Not all ACPI-devices have a HID + UID, allow specifying quirks for
    acpi_device_override_status() by path too.

    Note this moves the path/HID+UID check to after the CPU + DMI checks
    since the path lookup is somewhat costly.

    This way this lookup is only done on devices where the other checks
    match.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2dd4a54db4732849299ce4c0d4ee8938f7752f93
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon Nov 22 18:05:31 2021 +0100

    ACPI: Change acpi_device_always_present() into acpi_device_override_status()

    [ Upstream commit 1a68b346a2c9969c05e80a3b99a9ab160b5655c0 ]

    Currently, acpi_bus_get_status() calls acpi_device_always_present() to
    allow platform quirks to override the _STA return to report that a
    device is present (status = ACPI_STA_DEFAULT) independent of the _STA
    return.

    In some cases it might also be useful to have the opposite functionality
    and have a platform quirk which marks a device as not present (status = 0)
    to work around ACPI table bugs.

    Change acpi_device_always_present() into a more generic
    acpi_device_override_status() function to allow this.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a632c6b64e8c523a03a3e6c82df5d8626a11a82
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Mon Nov 22 18:05:30 2021 +0100

    ACPI / x86: Drop PWM2 device on Lenovo Yoga Book from always present table

    [ Upstream commit d431dfb764b145369be820fcdfd50f2159b9bbc2 ]

    It turns out that there is a WMI object which controls the PWM2 device
    used for the keyboard backlight and that WMI object also provides some
    other useful functionality.

    The upcoming lenovo-yogabook-wmi driver will offer both backlight
    control and the other functionality, so there no longer is a need
    to have the lpss-pwm driver binding to PWM2 for backlight control;
    and this is now actually undesirable because this will cause both
    the WMI code and the lpss-pwm driver to poke at the same PWM
    controller.

    Drop the always-present quirk for the PWM2 ACPI-device, so that the
     lpss-pwm controller will no longer bind to it.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 2963a60d750318ad6ae43d0f017eb5378987b0b7
Author: Zack Rusin <zackr@vmware.com>
Date:   Fri Nov 5 15:38:45 2021 -0400

    drm/vmwgfx: Introduce a new placement for MOB page tables

    [ Upstream commit f6be23264bbac88d1e2bb39658e1b8a397e3f46d ]

    For larger (bigger than a page) and noncontiguous mobs we have
    to create page tables that allow the host to find the memory.
    Those page tables just used regular system memory. Unfortunately
    in TTM those BO's are not allowed to be busy thus can't be
    fenced and we have to fence those bo's  because we don't want
    to destroy the page tables while the host is still executing
    the command buffers which might be accessing them.

    To solve it we introduce a new placement VMW_PL_SYSTEM which
    is very similar to TTM_PL_SYSTEM except that it allows
    fencing. This fixes kernel oops'es during unloading of the driver
    (and pci hot remove/add) which were caused by busy BO's in
    TTM_PL_SYSTEM being present in the delayed deletion list in
    TTM (TTM_PL_SYSTEM manager is destroyed before the delayed
    deletions are executed)

    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211105193845.258816-5-
zackr@vmware.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58d58bf87ace5d19865f202e6f27c547273a2f9d
Author: Zack Rusin <zackr@vmware.com>
Date:   Fri Nov 5 15:38:43 2021 -0400

    drm/vmwgfx: Release ttm memory if probe fails

    [ Upstream commit 28b5f3b6121b7db2a44be499cfca0b6b801588b6 ]

    The ttm mem global state was leaking if the vmwgfx driver load failed.

    In case of a driver load failure we have to make sure we also release
    the ttm mem global state.

    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211105193845.258816-3-
zackr@vmware.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 40f386c44a624be1b5066f81dbbbd8981c1f7810
Author: Adam Ward <Adam.Ward.opensource@diasemi.com>
Date:   Mon Nov 29 22:10:12 2021 +0000

    regulator: da9121: Prevent current limit change when enabled

    [ Upstream commit 24f0853228f3b98f1ef08d5824376c69bb8124d2 ]

    Prevent changing current limit when enabled as a precaution against
    possibile instability due to tight integration with switching cycle

    Signed-off-by: Adam Ward <Adam.Ward.opensource@diasemi.com>
    Link:
```

```
https://lore.kernel.org/r/52ee682476004a1736c1e0293358987319c1c415.1638223185.git.Adam.Ward.
opensource@diasemi.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fea9a6ed4e35bbc102e120ecfa071e4f8093c333
Author: Mansur Alisha Shaik <mansur@codeaurora.org>
Date:   Mon Nov 8 06:48:51 2021 +0100

    media: venus: avoid calling core_clk_setrate() concurrently during concurrent video
sessions

    [ Upstream commit 91f2b7d269e5c885c38c7ffa261f5276bd42f907 ]

    In existing implementation, core_clk_setrate() is getting called
    concurrently in concurrent video sessions. Before the previous call to
    core_clk_setrate returns, new call to core_clk_setrate is invoked from
    another video session running concurrently. This results in latest
    calculated frequency being set (higher/lower) instead of actual frequency
    required for that video session. It also results in stability crashes
    mention below. These resources are specific to video core, hence keeping
    under core lock would ensure that they are estimated for all running video
    sessions and called once for the video core.

    Crash logs:

    [    1.900089] WARNING: CPU: 4 PID: 1 at drivers/opp/debugfs.c:33
opp_debug_remove_one+0x2c/0x48
    [    1.908493] Modules linked in:
    [    1.911524] CPU: 4 PID: 1 Comm: swapper/0 Not tainted 5.10.67 #35
f8edb8c30cf2dd6838495dd9ef9be47af7f5f60c
    [    1.921036] Hardware name: Qualcomm Technologies, Inc. sc7280 IDP SKU2 platform (DT)
    [    1.928673] pstate: 60800009 (nZCv daif -PAN +UAO -TCO BTYPE=--)
    [    1.934608] pc : opp_debug_remove_one+0x2c/0x48
    [    1.939080] lr : opp_debug_remove_one+0x2c/0x48
    [    1.943560] sp : ffffffc011d7b7f0
    [    1.946836] pmr_save: 000000e0
    [    1.949854] x29: ffffffc011d7b7f0 x28: ffffffc010733bbc
    [    1.955104] x27: ffffffc010733ba8 x26: ffffff8083cedd00
    [    1.960355] x25: 0000000000000001 x24: 0000000000000000
    [    1.965603] x23: ffffff8083cc2878 x22: ffffff8083ceb900
    [    1.970852] x21: ffffff8083ceb910 x20: ffffff8083cc2800
    [    1.976101] x19: ffffff8083ceb900 x18: 00000000ffff0a10
    [    1.981352] x17: ffffff80837a5620 x16: 00000000000000ec
    [    1.986601] x15: ffffffc010519ad4 x14: 0000000000000003
    [    1.991849] x13: 0000000000000004 x12: 0000000000000001
    [    1.997100] x11: c0000000ffffdfff x10: 00000000ffffffff
    [    2.002348] x9 : d2627c580300dc00 x8 : d2627c580300dc00
    [    2.007596] x7 : 0720072007200720 x6 : ffffff80802ecf00
    [    2.012845] x5 : 0000000000190004 x4 : 0000000000000000
    [    2.018094] x3 : ffffffc011d7b478 x2 : ffffffc011d7b480
    [    2.023343] x1 : 00000000ffffdfff x0 : 0000000000000017
    [    2.028594] Call trace:
    [    2.031022]  opp_debug_remove_one+0x2c/0x48
    [    2.035160]  dev_pm_opp_put+0x94/0xb0
    [    2.038780]  _opp_remove_all+0x7c/0xc8
    [    2.042486]  _opp_remove_all_static+0x54/0x7c
    [    2.046796]  dev_pm_opp_remove_table+0x74/0x98
    [    2.051183]  devm_pm_opp_of_table_release+0x18/0x24
    [    2.056001]  devm_action_release+0x1c/0x28
    [    2.060053]  release_nodes+0x23c/0x2b8
    [    2.063760]  devres_release_group+0xcc/0xd0
    [    2.067900]  component_bind+0xac/0x168
    [    2.071608]  component_bind_all+0x98/0x124
    [    2.075664]  msm_drm_bind+0x1e8/0x678
```

```
      [    2.079287]  try_to_bring_up_master+0x60/0x134
      [    2.083674]  component_master_add_with_match+0xd8/0x120
      [    2.088834]  msm_pdev_probe+0x20c/0x2a0
      [    2.092629]  platform_drv_probe+0x9c/0xbc
      [    2.096598]  really_probe+0x11c/0x46c
      [    2.100217]  driver_probe_device+0x8c/0xf0
      [    2.104270]  device_driver_attach+0x54/0x78
      [    2.108407]  __driver_attach+0x48/0x148
      [    2.112201]  bus_for_each_dev+0x88/0xd4
      [    2.115998]  driver_attach+0x2c/0x38
      [    2.119534]  bus_add_driver+0x10c/0x200
      [    2.123330]  driver_register+0x6c/0x104
      [    2.127122]  __platform_driver_register+0x4c/0x58
      [    2.131767]  msm_drm_register+0x6c/0x70
      [    2.135560]  do_one_initcall+0x64/0x23c
      [    2.139357]  do_initcall_level+0xac/0x15c
      [    2.143321]  do_initcalls+0x5c/0x9c
      [    2.146778]  do_basic_setup+0x2c/0x38
      [    2.150401]  kernel_init_freeable+0xf8/0x15c
      [    2.154622]  kernel_init+0x1c/0x11c
      [    2.158079]  ret_from_fork+0x10/0x30
      [    2.161615] ---[ end trace a2cc45a0f784b212 ]---

      [    2.166272] Removing OPP: 300000000

      Signed-off-by: Mansur Alisha Shaik <mansur@codeaurora.org>
      Signed-off-by: Stanimir Varbanov <stanimir.varbanov@linaro.org>
      Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4ce72eb9f46016f7a27493d55d5de420c2697646
  Author: Sriram R <quic_srirrama@quicinc.com>
  Date:   Thu Nov 25 15:00:14 2021 +0530

      ath11k: Avoid NULL ptr access during mgmt tx cleanup

      [ Upstream commit a93789ae541c7d5c1c2a4942013adb6bcc5e2848 ]

      Currently 'ar' reference is not added in skb_cb during
      WMI mgmt tx. Though this is generally not used during tx completion
      callbacks, on interface removal the remaining idr cleanup callback
      uses the ar ptr from skb_cb from mgmt txmgmt_idr. Hence
      fill them during tx call for proper usage.

      Also free the skb which is missing currently in these
      callbacks.

      Crash_info:

      [19282.489476] Unable to handle kernel NULL pointer dereference at virtual address
  00000000
      [19282.489515] pgd = 91eb8000
      [19282.496702] [00000000] *pgd=00000000
      [19282.502524] Internal error: Oops: 5 [#1] PREEMPT SMP ARM
      [19282.783728] PC is at ath11k_mac_vif_txmgmt_idr_remove+0x28/0xd8 [ath11k]
      [19282.789170] LR is at idr_for_each+0xa0/0xc8

      Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-00729-QCAHKSWPL_SILICONZ-3 v2
      Signed-off-by: Sriram R <quic_srirrama@quicinc.com>
      Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
      Link: https://lore.kernel.org/r/1637832614-13831-1-git-send-email-
  quic_srirrama@quicinc.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit bd2b5268910cdd77a6b4376c4eee390d89f8e246
```

```
Author: Zekun Shen <bruceshenzk@gmail.com>
Date:   Fri Oct 29 16:19:23 2021 -0400

    rsi: Fix out-of-bounds read in rsi_read_pkt()

    [ Upstream commit f1cb3476e48b60c450ec3a1d7da0805bffc6e43a ]

    rsi_get_* functions rely on an offset variable from usb
    input. The size of usb input is RSI_MAX_RX_USB_PKT_SIZE(3000),
    while 2-byte offset can be up to 0xFFFF. Thus a large offset
    can cause out-of-bounds read.

    The patch adds a bound checking condition when rcv_pkt_len is 0,
    indicating it's USB. It's unclear whether this is triggerable
    from other type of bus. The following check might help in that case.
    offset > rcv_pkt_len - FRAME_DESC_SZ

    The bug is trigerrable with conpromised/malfunctioning USB devices.
    I tested the patch with the crashing input and got no more bug report.

    Attached is the KASAN report from fuzzing.

    BUG: KASAN: slab-out-of-bounds in rsi_read_pkt+0x42e/0x500 [rsi_91x]
    Read of size 2 at addr ffff888019439fdb by task RX-Thread/227

    CPU: 0 PID: 227 Comm: RX-Thread Not tainted 5.6.0 #66
    Call Trace:
     dump_stack+0x76/0xa0
     print_address_description.constprop.0+0x16/0x200
     ? rsi_read_pkt+0x42e/0x500 [rsi_91x]
     ? rsi_read_pkt+0x42e/0x500 [rsi_91x]
     __kasan_report.cold+0x37/0x7c
     ? rsi_read_pkt+0x42e/0x500 [rsi_91x]
     kasan_report+0xe/0x20
     rsi_read_pkt+0x42e/0x500 [rsi_91x]
     rsi_usb_rx_thread+0x1b1/0x2fc [rsi_usb]
     ? rsi_probe+0x16a0/0x16a0 [rsi_usb]
     ? _raw_spin_lock_irqsave+0x7b/0xd0
     ? _raw_spin_trylock_bh+0x120/0x120
     ? __wake_up_common+0x10b/0x520
     ? rsi_probe+0x16a0/0x16a0 [rsi_usb]
     kthread+0x2b5/0x3b0
     ? kthread_create_on_node+0xd0/0xd0
     ret_from_fork+0x22/0x40

    Reported-by: Brendan Dolan-Gavitt <brendandg@nyu.edu>
    Signed-off-by: Zekun Shen <bruceshenzk@gmail.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/YXxXS4wgu2OsmlVv@10-18-43-117.dynapool.wireless.nyu.edu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b0c55818d1514756e624911704b808b1a37485b
Author: Zekun Shen <bruceshenzk@gmail.com>
Date:   Fri Oct 29 15:49:03 2021 -0400

    rsi: Fix use-after-free in rsi_rx_done_handler()

    [ Upstream commit b07e3c6ebc0c20c772c0f54042e430acec2945c3 ]

    When freeing rx_cb->rx_skb, the pointer is not set to NULL,
    a later rsi_rx_done_handler call will try to read the freed
    address.
    This bug will very likley lead to double free, although
    detected early as use-after-free bug.
```

```
    The bug is triggerable with a compromised/malfunctional usb
    device. After applying the patch, the same input no longer
    triggers the use-after-free.

    Attached is the kasan report from fuzzing.

    BUG: KASAN: use-after-free in rsi_rx_done_handler+0x354/0x430 [rsi_usb]
    Read of size 4 at addr ffff8880188e5930 by task modprobe/231
    Call Trace:
     <IRQ>
     dump_stack+0x76/0xa0
     print_address_description.constprop.0+0x16/0x200
     ? rsi_rx_done_handler+0x354/0x430 [rsi_usb]
     ? rsi_rx_done_handler+0x354/0x430 [rsi_usb]
     __kasan_report.cold+0x37/0x7c
     ? dma_direct_unmap_page+0x90/0x110
     ? rsi_rx_done_handler+0x354/0x430 [rsi_usb]
     kasan_report+0xe/0x20
     rsi_rx_done_handler+0x354/0x430 [rsi_usb]
     __usb_hcd_giveback_urb+0x1e4/0x380
     usb_giveback_urb_bh+0x241/0x4f0
     ? __usb_hcd_giveback_urb+0x380/0x380
     ? apic_timer_interrupt+0xa/0x20
     tasklet_action_common.isra.0+0x135/0x330
     __do_softirq+0x18c/0x634
     ? handle_irq_event+0xcd/0x157
     ? handle_edge_irq+0x1eb/0x7b0
     irq_exit+0x114/0x140
     do_IRQ+0x91/0x1e0
     common_interrupt+0xf/0xf
     </IRQ>

    Reported-by: Brendan Dolan-Gavitt <brendandg@nyu.edu>
    Signed-off-by: Zekun Shen <bruceshenzk@gmail.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/YXxQL/vIiYcZUu/j@10-18-43-117.dynapool.wireless.nyu.edu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d3989c5050f10ae9bbec9f32492b500420d04a1
Author: Zekun Shen <bruceshenzk@gmail.com>
Date:   Sat Oct 30 22:42:50 2021 -0400

    mwifiex: Fix skb_over_panic in mwifiex_usb_recv()

    [ Upstream commit 04d80663f67ccef893061b49ec8a42ff7045ae84 ]

    Currently, with an unknown recv_type, mwifiex_usb_recv
    just return -1 without restoring the skb. Next time
    mwifiex_usb_rx_complete is invoked with the same skb,
    calling skb_put causes skb_over_panic.

    The bug is triggerable with a compromised/malfunctioning
    usb device. After applying the patch, skb_over_panic
    no longer shows up with the same input.

    Attached is the panic report from fuzzing.
    skbuff: skb_over_panic: text:000000003bf1b5fa
     len:2048 put:4 head:00000000dd6a115b data:000000000a9445d8
     tail:0x844 end:0x840 dev:<NULL>
    kernel BUG at net/core/skbuff.c:109!
    invalid opcode: 0000 [#1] SMP KASAN NOPTI
    CPU: 0 PID: 198 Comm: in:imklog Not tainted 5.6.0 #60
    RIP: 0010:skb_panic+0x15f/0x161
    Call Trace:
     <IRQ>
```

```
       ? mwifiex_usb_rx_complete+0x26b/0xfcd [mwifiex_usb]
       skb_put.cold+0x24/0x24
       mwifiex_usb_rx_complete+0x26b/0xfcd [mwifiex_usb]
       __usb_hcd_giveback_urb+0x1e4/0x380
       usb_giveback_urb_bh+0x241/0x4f0
       ? __hrtimer_run_queues+0x316/0x740
       ? __usb_hcd_giveback_urb+0x380/0x380
       tasklet_action_common.isra.0+0x135/0x330
       __do_softirq+0x18c/0x634
       irq_exit+0x114/0x140
       smp_apic_timer_interrupt+0xde/0x380
       apic_timer_interrupt+0xf/0x20
       </IRQ>

      Reported-by: Brendan Dolan-Gavitt <brendandg@nyu.edu>
      Signed-off-by: Zekun Shen <bruceshenzk@gmail.com>
      Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
      Link: https://lore.kernel.org/r/YX4CqjfRcTa6bVL+@Zekuns-MBP-16.fios-router.home
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b888433c205225ee0c8debea55a6cfbe1413763b
  Author: Ping-Ke Shih <pkshih@realtek.com>
  Date:   Fri Nov 19 13:57:29 2021 +0800

      rtw89: fix potentially access out of range of RF register array

      [ Upstream commit 30101812a09b37bc8aa409a83f603d4c072198f2 ]

      The RF register array is used to help firmware to restore RF settings.
      The original code can potentially access out of range, if the size is
      between (RTW89_H2C_RF_PAGE_SIZE * RTW89_H2C_RF_PAGE_NUM + 1) to
      ((RTW89_H2C_RF_PAGE_SIZE + 1) * RTW89_H2C_RF_PAGE_NUM). Fortunately,
      current used size doesn't fall into the wrong case, and the size will not
      change if we don't update RF parameter.

      Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
      Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
      Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
      Link: https://lore.kernel.org/r/20211119055729.12826-1-pkshih@realtek.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 91b10f3fcffe17b62da72795666164b5c15a0389
  Author: Stephan Müller <smueller@chronox.de>
  Date:   Sun Nov 21 15:14:20 2021 +0100

      crypto: jitter - consider 32 LSB for APT

      [ Upstream commit 552d03a223eda3df84526ab2c1f4d82e15eaee7a ]

      The APT compares the current time stamp with a pre-set value. The
      current code only considered the 4 LSB only. Yet, after reviews by
      mathematicians of the user space Jitter RNG version >= 3.1.0, it was
      concluded that the APT can be calculated on the 32 LSB of the time
      delta. Thi change is applied to the kernel.

      This fixes a bug where an AMD EPYC fails this test as its RDTSC value
      contains zeros in the LSB. The most appropriate fix would have been to
      apply a GCD calculation and divide the time stamp by the GCD. Yet, this
      is a significant code change that will be considered for a future
      update. Note, tests showed that constantly the GCD always was 32 on
      these systems, i.e. the 5 LSB were always zero (thus failing the APT
      since it only considered the 4 LSB for its calculation).

      Signed-off-by: Stephan Mueller <smueller@chronox.de>
      Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
```

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

commit aefcf6a3db81eded035c1242e300f25dde9a0781
Author: Chengfeng Ye <cyeaa@connect.ust.hk>
Date:   Fri Nov 5 06:45:07 2021 -0700

    HSI: core: Fix return freed object in hsi_new_client

    [ Upstream commit a1ee1c08fcd5af03187dcd41dcab12fd5b379555 ]

    cl is freed on error of calling device_register, but this
    object is return later, which will cause uaf issue. Fix it
    by return NULL on error.

    Signed-off-by: Chengfeng Ye <cyeaa@connect.ust.hk>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d42c29d28db82b489ad3f81d50d45f7281d3a48
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Nov 25 21:30:10 2021 +0100

    gpiolib: acpi: Do not set the IRQ type if the IRQ is already in use

    [ Upstream commit bdfd6ab8fdccd8b138837efff66f4a1911496378 ]

    If the IRQ is already in use, then acpi_dev_gpio_irq_get_by() really
    should not change the type underneath the current owner.

    I specifically hit an issue with this an a Chuwi Hi8 Super (CWI509) Bay
    Trail tablet, when the Boot OS selection in the BIOS is set to Android.
    In this case _STA for a MAX17047 ACPI I2C device wrongly returns 0xf and
    the _CRS resources for this device include a GpioInt pointing to a GPIO
    already in use by an _AEI handler, with a different type then specified
    in the _CRS for the MAX17047 device. Leading to the acpi_dev_gpio_irq_get()
    call done by the i2c-core-acpi.c code changing the type breaking the
    _AEI handler.

    Now this clearly is a bug in the DSDT of this tablet (in Android mode),
    but in general calling irq_set_irq_type() on an IRQ which already is
    in use seems like a bad idea.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1f6d3f6c99666d8eed890e7c95ecb87ec2a1c93
Author: Fugang Duan <fugang.duan@nxp.com>
Date:   Thu Nov 25 10:03:49 2021 +0800

    tty: serial: imx: disable UCR4_OREN in .stop_rx() instead of .shutdown()

    [ Upstream commit 028e083832b06fdeeb290e1e57dc1f6702c4c215 ]

    The UCR4_OREN should be disabled before disabling the uart receiver in
    .stop_rx() instead of in the .shutdown().

    Otherwise, if we have the overrun error during the receiver disable
    process, the overrun interrupt will keep trigging until we disable the
    OREN interrupt in the .shutdown(), because the ORE status can only be
    cleared when read the rx FIFO or reset the controller.  Although the
    called time between the receiver disable and OREN disable in .shutdown()
    is very short, there is still the risk of endless interrupt during this
    short period of time. So here change to disable OREN before the receiver
    been disabled in .stop_rx().
```

```
    Signed-off-by: Fugang Duan <fugang.duan@nxp.com>
    Signed-off-by: Sherry Sun <sherry.sun@nxp.com>
    Link: https://lore.kernel.org/r/20211125020349.4980-1-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35884eb8f7064b704ff87fb05a1fa0add2faa0b3
Author: Emil Renner Berthing <kernel@esmil.dk>
Date:   Tue Nov 16 16:01:17 2021 +0100

    serial: 8250_dw: Add StarFive JH7100 quirk

    [ Upstream commit 57dcb6ec85d59e04285b7dcf10924bb819c8e46f ]

    On the StarFive JH7100 RISC-V SoC the UART core clocks can't be set to
    exactly 16 * 115200Hz and many other common bitrates. Trying this will
    only result in a higher input clock, but low enough that the UART's
    internal divisor can't come close enough to the baud rate target.
    So rather than try to set the input clock it's better to skip the
    clk_set_rate call and rely solely on the UART's internal divisor.

    Signed-off-by: Emil Renner Berthing <kernel@esmil.dk>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Reviewed-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Link: https://lore.kernel.org/r/20211116150119.2171-15-kernel@esmil.dk
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ac1bcc94bfbcc461c49d124447d2e4fd113608a
Author: Jiri Slaby <jirislaby@kernel.org>
Date:   Thu Nov 18 08:31:22 2021 +0100

    mxser: increase buf_overrun if tty_insert_flip_char() fails

    [ Upstream commit eb68ac0462bffc2ceb63b3a76737d6c9f186e6de ]

    mxser doesn't increase port->icount.buf_overrun at all. Do so if overrun
    happens, so that it can be read from the stats.

    Signed-off-by: Jiri Slaby <jslaby@suse.cz>
    Link: https://lore.kernel.org/r/20211118073125.12283-17-jslaby@suse.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit beffe85935f3ea3e71d7cd4a1fa3bab77726b770
Author: Jiri Slaby <jirislaby@kernel.org>
Date:   Thu Nov 18 08:31:20 2021 +0100

    mxser: don't throttle manually

    [ Upstream commit c6693e6e07805f1b7822b13a5b482bf2b6a1f312 ]

    First, checking tty->receive_room to signalize whether there is enough space
    in the tty buffers does not make much sense. Provided the tty buffers
    are in tty_port and those are not checked at all.

    Second, if the rx path is throttled, with CRTSCTS, RTS is deasserted,
    but is never asserted again. This leads to port "lockup", not accepting
    any more input.

    So:
    1) stty -F /dev/ttyMI0 crtscts # the mxser port
    2) stty -F /dev/ttyS6 crtscts # the connected port
    3) cat /dev/ttyMI0
```

```
    4) "write in a loop" to /dev/ttyS6
    5) cat from 3) produces the bytes from 4)
    6) killall -STOP cat (the 3)'s one)
    7) wait for RTS to drop on /dev/ttyMI0
    8) killall -CONT cat (again the 3)'s one)

    cat erroneously produces no more output now (i.e. no data sent from
    ttyS6 to ttyMI can be seen).

    Note that the step 7) is performed twice: once from n_tty by
    tty_throttle_safe(), once by mxser_stoprx() from the receive path. Then
    after step 7), n_tty correctly unthrottles the input, but mxser calls
    mxser_stoprx() again as there is still only a little space in n_tty
    buffers (tty->receive_room mentioned at the beginning), but the device's
    FIFO is/can be already filled.

    After this patch, the output is correctly resumed, i.e. n_tty both
    throttles and unthrottles without interfering with mxser's attempts.

    This allows us to get rid of the non-standard ldisc_stop_rx flag from
    struct mxser_port.

    Signed-off-by: Jiri Slaby <jslaby@suse.cz>
    Link: https://lore.kernel.org/r/20211118073125.12283-15-jslaby@suse.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 724eeafcc1fb8f0c770646c48b4a0cc63b339cc8
Author: Jiri Slaby <jirislaby@kernel.org>
Date:   Thu Nov 18 08:31:09 2021 +0100

    mxser: keep only !tty test in ISR

    [ Upstream commit 274ab58dc2b460cc474ffc7ccfcede4b2be1a3f5 ]

    The others are superfluous with tty refcounting in place now. And they
    are racy in fact:
    * tty_port_initialized() reports false for a small moment after
      interrupts are enabled.
    * closing is 1 while the port is still alive.

    The queues are flushed later during close anyway. So there is no need
    for this special handling. Actually, the ISR should not flush the
    queues. It should behave as every other driver, just queue the chars
    into tty buffer and go on. But this will be changed later. There is
    still a lot code depending on having tty in ISR (and not only tty_port).

    Signed-off-by: Jiri Slaby <jslaby@suse.cz>
    Link: https://lore.kernel.org/r/20211118073125.12283-4-jslaby@suse.cz
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e449461989f0674f188d30494ffd3f5ba74c1a05
Author: Martyn Welch <martyn.welch@collabora.com>
Date:   Thu Nov 25 10:53:02 2021 +0000

    drm/bridge: megachips: Ensure both bridges are probed before registration

    [ Upstream commit 11632d4aa2b3f126790e81a4415d6c23103cf8bb ]

    In the configuration used by the b850v3, the STDP2690 is used to read EDID
    data whilst it's the STDP4028 which can detect when monitors are connected.

    This can result in problems at boot with monitors connected when the
    STDP4028 is probed first, a monitor is detected and an attempt is made to
```

```
    read the EDID data before the STDP2690 has probed:

    [    3.795721] Unable to handle kernel NULL pointer dereference at virtual address
 00000018
    [    3.803845] pgd = (ptrval)
    [    3.806581] [00000018] *pgd=00000000
    [    3.810180] Internal error: Oops: 5 [#1] SMP ARM
    [    3.814813] Modules linked in:
    [    3.817879] CPU: 0 PID: 64 Comm: kworker/u4:1 Not tainted 5.15.0 #1
    [    3.824161] Hardware name: Freescale i.MX6 Quad/DualLite (Device Tree)
    [    3.830705] Workqueue: events_unbound deferred_probe_work_func
    [    3.836565] PC is at stdp2690_get_edid+0x44/0x19c
    [    3.841286] LR is at ge_b850v3_lvds_get_modes+0x2c/0x5c
    [    3.846526] pc : [<805eae10>]    lr : [<805eb138>]    psr: 80000013
    [    3.852802] sp : 81c359d0  ip : 7dbb550b  fp : 81c35a1c
    [    3.858037] r10: 81c73840  r9 : 81c73894  r8 : 816d9800
    [    3.863270] r7 : 00000000  r6 : 81c34000  r5 : 00000000  r4 : 810c35f0
    [    3.869808] r3 : 80e3e294  r2 : 00000080  r1 : 00000cc0  r0 : 81401180
    [    3.876349] Flags: Nzcv  IRQs on  FIQs on  Mode SVC_32  ISA ARM  Segment none
    [    3.883499] Control: 10c5387d  Table: 1000404a  DAC: 00000051
    [    3.889254] Register r0 information: slab kmem_cache start 81401180 pointer offset 0
    [    3.897034] Register r1 information: non-paged memory
    [    3.902097] Register r2 information: non-paged memory
    [    3.907160] Register r3 information: non-slab/vmalloc memory
    [    3.912832] Register r4 information: non-slab/vmalloc memory
    [    3.918503] Register r5 information: NULL pointer
    [    3.923217] Register r6 information: non-slab/vmalloc memory
    [    3.928887] Register r7 information: NULL pointer
    [    3.933601] Register r8 information: slab kmalloc-1k start 816d9800 pointer offset 0
 size 1024
    [    3.942244] Register r9 information: slab kmalloc-2k start 81c73800 pointer offset
 148 size 2048
    [    3.951058] Register r10 information: slab kmalloc-2k start 81c73800 pointer offset
 64 size 2048
    [    3.959873] Register r11 information: non-slab/vmalloc memory
    [    3.965632] Register r12 information: non-paged memory
    [    3.970781] Process kworker/u4:1 (pid: 64, stack limit = 0x(ptrval))
    [    3.977148] Stack: (0x81c359d0 to 0x81c36000)
    [    3.981517] 59c0:                                     80b2b668 80b2b5bc 000002e2
 0000034e
    [    3.989712] 59e0: 81c35a8c 816d98e8 81c35a14 7dbb550b 805bfcd0 810c35f0 81c73840
 824addc0
    [    3.997906] 5a00: 00001000 816d9800 81c73894 81c73840 81c35a34 81c35a20 805eb138
 805eadd8
    [    4.006099] 5a20: 810c35f0 00000045 81c35adc 81c35a38 80594188 805eb118 80d7c788
 80dd1848
    [    4.014292] 5a40: 00000000 81c35a50 80dca950 811194d3 80dca7c4 80dca944 80dca91c
 816d9800
    [    4.022485] 5a60: 81c34000 81c760a8 816d9800 80c58c98 810c35f0 816d98e8 00001000
 00001000
    [    4.030678] 5a80: 00000000 00000000 8017712c 81c60000 00000002 00000001 00000000
 00000000
    [    4.038870] 5aa0: 816d9900 816d9900 00000000 7dbb550b 805c700c 00000008 826282c8
 826282c8
    [    4.047062] 5ac0: 00001000 81e1ce40 00001000 00000002 81c35bf4 81c35ae0 805d9694
 80593fc0
    [    4.055255] 5ae0: 8017a970 80179ad8 00000179 00000000 81c35bcc 81c35b00 80177108
 8017a950
    [    4.063447] 5b00: 00000000 81c35b10 81c34000 00000000 81004fd8 81010a38 00000000
 00000059
    [    4.071639] 5b20: 816d98d4 81fbb718 00000013 826282c8 8017a940 81c35b40 81134448
 00000400
    [    4.079831] 5b40: 00000178 00000000 e063b9c1 00000000 c2000049 00000040 00000000
 00000008
    [    4.088024] 5b60: 82628300 82628380 00000000 00000000 81c34000 00000000 81fbb700
```

```
82628340
    [    4.096216] 5b80: 826283c0 00001000 00000000 00000010 816d9800 826282c0 801766f8
00000000
    [    4.104408] 5ba0: 00000000 81004fd8 00000049 00000000 00000000 00000001 80dcf940
80178de4
    [    4.112601] 5bc0: 81c35c0c 7dbb550b 80178de4 81fbb700 00000010 00000010 810c35f4
81e1ce40
    [    4.120793] 5be0: 81c40908 0000000c 81c35c64 81c35bf8 805a7f18 805d94a0 81c35c3c
816d9800
    [    4.128985] 5c00: 00000010 81c34000 81c35c2c 81c35c18 8012fce0 805be90c 81c35c3c
81c35c28
    [    4.137178] 5c20: 805be90c 80173210 81fbb600 81fbb6b4 81c35c5c 7dbb550b 81c35c64
81fbb700
    [    4.145370] 5c40: 816d9800 00000010 810c35f4 81e1ce40 81c40908 0000000c 81c35c84
81c35c68
    [    4.153565] 5c60: 805a8c78 805a7ed0 816d9800 81fbb700 00000010 00000000 81c35cac
81c35c88
    [    4.161758] 5c80: 805a8dc4 805a8b68 816d9800 00000000 816d9800 00000000 8179f810
810c42d0
    [    4.169950] 5ca0: 81c35ccc 81c35cb0 805e47b0 805a8d18 824aa240 81e1ea80 81c40908
81126b60
    [    4.178144] 5cc0: 81c35d14 81c35cd0 8060db1c 805e46cc 81c35d14 81c35ce0 80dd90f8
810c4d58
    [    4.186338] 5ce0: 80dd90dc 81fe9740 fffffffe 81fe9740 81e1ea80 00000000 810c4d6c
80c4b95c
    [    4.194531] 5d00: 80dd9a3c 815c6810 81c35d34 81c35d18 8060dc9c 8060d8fc 8246b440
815c6800
    [    4.202724] 5d20: 815c6810 eefd8e00 81c35d44 81c35d38 8060dd80 8060dbec 81c35d6c
81c35d48
    [    4.210918] 5d40: 805e98a4 8060dd70 00000000 815c6810 810c45b0 81126e90 81126e90
80dd9a3c
    [    4.219112] 5d60: 81c35d8c 81c35d70 80619574 805e9808 815c6810 00000000 810c45b0
81126e90
    [    4.227305] 5d80: 81c35db4 81c35d90 806168dc 80619514 80625df0 80623c80 815c6810
810c45b0
    [    4.235498] 5da0: 81c35e6c 815c6810 81c35dec 81c35db8 80616d04 80616800 81c35de4
81c35dc8
    [    4.243691] 5dc0: 808382b0 80b2f444 8116e310 8116e314 81c35e6c 815c6810 00000003
80dd9a3c
    [    4.251884] 5de0: 81c35e14 81c35df0 80616ec8 80616c60 00000001 810c45b0 81c35e6c
815c6810
    [    4.260076] 5e00: 00000001 80dd9a3c 81c35e34 81c35e18 80617338 80616e90 00000000
81c35e6c
    [    4.268269] 5e20: 80617284 81c34000 81c35e64 81c35e38 80614730 80617290 81c35e64
8171a06c
    [    4.276461] 5e40: 81e220b8 7dbb550b 815c6810 81c34000 815c6854 81126e90 81c35e9c
81c35e68
    [    4.284654] 5e60: 8061673c 806146a8 8060f5e0 815c6810 00000001 7dbb550b 00000000
810c5080
    [    4.292847] 5e80: 810c5320 815c6810 81126e90 00000000 81c35eac 81c35ea0 80617554
80616650
    [    4.301040] 5ea0: 81c35ecc 81c35eb0 80615694 80617544 810c5080 810c5080 810c5094
81126e90
    [    4.309233] 5ec0: 81c35efc 81c35ed0 80615c6c 8061560c 80615bc0 810c50c0 817eeb00
81412800
    [    4.317425] 5ee0: 814c3000 00000000 814c300d 81119a60 81c35f3c 81c35f00 80141488
80615bcc
    [    4.325618] 5f00: 81c60000 81c34000 81c35f24 81c35f18 80143078 817eeb00 81412800
817eeb18
    [    4.333811] 5f20: 81412818 81003d00 00000088 81412800 81c35f74 81c35f40 80141a48
80141298
    [    4.342005] 5f40: 81c35f74 81c34000 801481ac 817efa40 817efc00 801417d8 817eeb00
00000000
    [    4.350199] 5f60: 815a7e7c 81c34000 81c35fac 81c35f78 80149b1c 801417e4 817efc20
817efc20
```

```
     [    4.358391] 5f80: ffffe000 817efa40 801499a8 00000000 00000000 00000000 00000000
  00000000
     [    4.366583] 5fa0: 00000000 81c35fb0 80100130 801499b4 00000000 00000000 00000000
  00000000
     [    4.374774] 5fc0: 00000000 00000000 00000000 00000000 00000000 00000000 00000000
  00000000
     [    4.382966] 5fe0: 00000000 00000000 00000000 00000000 00000013 00000000 00000000
  00000000
     [    4.391155] Backtrace:
     [    4.393613] [<805eadcc>] (stdp2690_get_edid) from [<805eb138>]
  (ge_b850v3_lvds_get_modes+0x2c/0x5c)
     [    4.402691]  r10:81c73840 r9:81c73894 r8:816d9800 r7:00001000 r6:824addc0 r5:81c73840
     [    4.410534]  r4:810c35f0
     [    4.413073] [<805eb10c>] (ge_b850v3_lvds_get_modes) from [<80594188>]
  (drm_helper_probe_single_connector_modes+0x1d4/0x84c)
     [    4.424240]  r5:00000045 r4:810c35f0
     [    4.427822] [<80593fb4>] (drm_helper_probe_single_connector_modes) from [<805d9694>]
  (drm_client_modeset_probe+0x200/0x1384)
     [    4.439074]  r10:00000002 r9:00001000 r8:81e1ce40 r7:00001000 r6:826282c8 r5:826282c8
     [    4.446917]  r4:00000008
     [    4.449455] [<805d9494>] (drm_client_modeset_probe) from [<805a7f18>]
  (__drm_fb_helper_initial_config_and_unlock+0x54/0x5b4)
     [    4.460713]  r10:0000000c r9:81c40908 r8:81e1ce40 r7:810c35f4 r6:00000010 r5:00000010
     [    4.468556]  r4:81fbb700
     [    4.471095] [<805a7ec4>] (__drm_fb_helper_initial_config_and_unlock) from
  [<805a8c78>] (drm_fbdev_client_hotplug+0x11c/0x1b0)
     [    4.482434]  r10:0000000c r9:81c40908 r8:81e1ce40 r7:810c35f4 r6:00000010 r5:816d9800
     [    4.490276]  r4:81fbb700
     [    4.492814] [<805a8b5c>] (drm_fbdev_client_hotplug) from [<805a8dc4>]
  (drm_fbdev_generic_setup+0xb8/0x1a4)
     [    4.502494]  r7:00000000 r6:00000010 r5:81fbb700 r4:816d9800
     [    4.508160] [<805a8d0c>] (drm_fbdev_generic_setup) from [<805e47b0>]
  (imx_drm_bind+0xf0/0x130)
     [    4.516805]  r7:810c42d0 r6:8179f810 r5:00000000 r4:816d9800
     [    4.522474] [<805e46c0>] (imx_drm_bind) from [<8060db1c>]
  (try_to_bring_up_master+0x22c/0x2f0)
     [    4.531116]  r7:81126b60 r6:81c40908 r5:81e1ea80 r4:824aa240
     [    4.536783] [<8060d8f0>] (try_to_bring_up_master) from [<8060dc9c>]
  (__component_add+0xbc/0x184)
     [    4.545597]  r10:815c6810 r9:80dd9a3c r8:80c4b95c r7:810c4d6c r6:00000000 r5:81e1ea80
     [    4.553440]  r4:81fe9740
     [    4.555980] [<8060dbe0>] (__component_add) from [<8060dd80>]
  (component_add+0x1c/0x20)
     [    4.563921]  r7:eefd8e00 r6:815c6810 r5:815c6800 r4:8246b440
     [    4.569589] [<8060dd64>] (component_add) from [<805e98a4>]
  (dw_hdmi_imx_probe+0xa8/0xe8)
     [    4.577702] [<805e97fc>] (dw_hdmi_imx_probe) from [<80619574>]
  (platform_probe+0x6c/0xc8)
     [    4.585908]  r9:80dd9a3c r8:81126e90 r7:81126e90 r6:810c45b0 r5:815c6810 r4:00000000
     [    4.593662] [<80619508>] (platform_probe) from [<806168dc>] (really_probe+0xe8/0x460)
     [    4.601524]  r7:81126e90 r6:810c45b0 r5:00000000 r4:815c6810
     [    4.607191] [<806167f4>] (really_probe) from [<80616d04>]
  (__driver_probe_device+0xb0/0x230)
     [    4.615658]  r7:815c6810 r6:81c35e6c r5:810c45b0 r4:815c6810
     [    4.621326] [<80616c54>] (__driver_probe_device) from [<80616ec8>]
  (driver_probe_device+0x44/0xe0)
     [    4.630313]  r9:80dd9a3c r8:00000003 r7:815c6810 r6:81c35e6c r5:8116e314 r4:8116e310
     [    4.638068] [<80616e84>] (driver_probe_device) from [<80617338>]
  (__device_attach_driver+0xb4/0x12c)
     [    4.647227]  r9:80dd9a3c r8:00000001 r7:815c6810 r6:81c35e6c r5:810c45b0 r4:00000001
     [    4.654981] [<80617284>] (__device_attach_driver) from [<80614730>]
  (bus_for_each_drv+0x94/0xd8)
     [    4.663794]  r7:81c34000 r6:80617284 r5:81c35e6c r4:00000000
     [    4.669461] [<8061469c>] (bus_for_each_drv) from [<8061673c>]
  (__device_attach+0xf8/0x190)
```

```
      [    4.677753]  r7:81126e90 r6:815c6854 r5:81c34000 r4:815c6810
      [    4.683419] [<80616644>] (__device_attach) from [<80617554>]
  (device_initial_probe+0x1c/0x20)
      [    4.691971]  r8:00000000 r7:81126e90 r6:815c6810 r5:810c5320 r4:810c5080
      [    4.698681] [<80617538>] (device_initial_probe) from [<80615694>]
  (bus_probe_device+0x94/0x9c)
      [    4.707318] [<80615600>] (bus_probe_device) from [<80615c6c>]
  (deferred_probe_work_func+0xac/0xf0)
      [    4.716305]  r7:81126e90 r6:810c5094 r5:810c5080 r4:810c5080
      [    4.721973] [<80615bc0>] (deferred_probe_work_func) from [<80141488>]
  (process_one_work+0x1fc/0x54c)
      [    4.731139]  r10:81119a60 r9:814c300d r8:00000000 r7:814c3000 r6:81412800 r5:817eeb00
      [    4.738981]  r4:810c50c0 r3:80615bc0
      [    4.742563] [<8014128c>] (process_one_work) from [<80141a48>]
  (worker_thread+0x270/0x570)
      [    4.750765]  r10:81412800 r9:00000088 r8:81003d00 r7:81412818 r6:817eeb18 r5:81412800
      [    4.758608]  r4:817eeb00
      [    4.761147] [<801417d8>] (worker_thread) from [<80149b1c>] (kthread+0x174/0x190)
      [    4.768574]  r10:81c34000 r9:815a7e7c r8:00000000 r7:817eeb00 r6:801417d8 r5:817efc00
      [    4.776417]  r4:817efa40
      [    4.778955] [<801499a8>] (kthread) from [<80100130>] (ret_from_fork+0x14/0x24)
      [    4.786201] Exception stack(0x81c35fb0 to 0x81c35ff8)
      [    4.791266] 5fa0:                                     00000000 00000000 00000000
  00000000
      [    4.799459] 5fc0: 00000000 00000000 00000000 00000000 00000000 00000000 00000000
  00000000
      [    4.807651] 5fe0: 00000000 00000000 00000000 00000000 00000013 00000000
      [    4.814279]  r10:00000000 r9:00000000 r8:00000000 r7:00000000 r6:00000000 r5:801499a8
      [    4.822120]  r4:817efa40
      [    4.824664] Code: e3a02080 e593001c e3a01d33 e3a05000 (e5979018)

      Split the registration from the STDP4028 probe routine and only perform
      registration once both the STDP4028 and STDP2690 have probed.

      Signed-off-by: Martyn Welch <martyn.welch@collabora.com>
      CC: Peter Senna Tschudin <peter.senna@gmail.com>
      CC: Martyn Welch <martyn.welch@collabora.co.uk>
      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@gmail.com>
      Signed-off-by: Robert Foss <robert.foss@linaro.org>
      Link:
  https://patchwork.freedesktop.org/patch/msgid/43552c3404e8fdf92d8bc5658fac24e9f03c2c57.16378
  36606.git.martyn.welch@collabora.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c706b4cde72c9ae8278054a22aa390a339b99c4c
  Author: Martin Leung <Martin.Leung@amd.com>
  Date:   Fri Nov 12 17:59:31 2021 -0500

      drm/amd/display: add else to avoid double destroy clk_mgr

      [ Upstream commit 11dff0e871037a6ad978e52f826a2eb7f5fb274a ]

      [Why & How]
      when changing some code we accidentally
      changed else if-> if. reverting that.

      Reviewed-by: Aric Cyr <Aric.Cyr@amd.com>
      Acked-by: Qingqing Zhuo <qingqing.zhuo@amd.com>
      Signed-off-by: Martin Leung <Martin.Leung@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 5cccdf527a1bc6bc9c2a6c1d36a3b9fdb232f152
Author: Danielle Ratson <danieller@nvidia.com>
Date:   Tue Nov 23 09:54:47 2021 +0200

    mlxsw: pci: Add shutdown method in PCI driver

    [ Upstream commit c1020d3cf4752f61a6a413f632ea2ce2370e150d ]

    On an arm64 platform with the Spectrum ASIC, after loading and executing
    a new kernel via kexec, the following trace [1] is observed. This seems
    to be caused by the fact that the device is not properly shutdown before
    executing the new kernel.

    Fix this by implementing a shutdown method which mirrors the remove
    method, as recommended by the kexec maintainer [2][3].

    [1]
    BUG: Bad page state in process devlink pfn:22f73d
    page:fffffe00089dcf40 refcount:-1 mapcount:0 mapping:0000000000000000 index:0x0
    flags: 0x2ffff00000000000()
    raw: 2ffff00000000000 0000000000000000 ffffffff089d0201 0000000000000000
    raw: 0000000000000000 0000000000000000 ffffffffffffffff 0000000000000000
    page dumped because: nonzero _refcount
    Modules linked in:
    CPU: 1 PID: 16346 Comm: devlink Tainted: G B 5.8.0-rc6-custom-273020-gac6b365b1bf5 #44
    Hardware name: Marvell Armada 7040 TX4810M (DT)
    Call trace:
     dump_backtrace+0x0/0x1d0
     show_stack+0x1c/0x28
     dump_stack+0xbc/0x118
     bad_page+0xcc/0xf8
     check_free_page_bad+0x80/0x88
     __free_pages_ok+0x3f8/0x418
     __free_pages+0x38/0x60
     kmem_freepages+0x200/0x2a8
     slab_destroy+0x28/0x68
     slabs_destroy+0x60/0x90
     ___cache_free+0x1b4/0x358
     kfree+0xc0/0x1d0
     skb_free_head+0x2c/0x38
     skb_release_data+0x110/0x1a0
     skb_release_all+0x2c/0x38
     consume_skb+0x38/0x130
     __dev_kfree_skb_any+0x44/0x50
     mlxsw_pci_rdq_fini+0x8c/0xb0
     mlxsw_pci_queue_fini.isra.0+0x28/0x58
     mlxsw_pci_queue_group_fini+0x58/0x88
     mlxsw_pci_aqs_fini+0x2c/0x60
     mlxsw_pci_fini+0x34/0x50
     mlxsw_core_bus_device_unregister+0x104/0x1d0
     mlxsw_devlink_core_bus_device_reload_down+0x2c/0x48
     devlink_reload+0x44/0x158
     devlink_nl_cmd_reload+0x270/0x290
     genl_rcv_msg+0x188/0x2f0
     netlink_rcv_skb+0x5c/0x118
     genl_rcv+0x3c/0x50
     netlink_unicast+0x1bc/0x278
     netlink_sendmsg+0x194/0x390
     __sys_sendto+0xe0/0x158
     __arm64_sys_sendto+0x2c/0x38
     el0_svc_common.constprop.0+0x70/0x168
     do_el0_svc+0x28/0x88
     el0_sync_handler+0x88/0x190
```

```
      el0_sync+0x140/0x180

    [2]
    https://www.mail-archive.com/linux-kernel@vger.kernel.org/msg1195432.html

    [3]
    https://patchwork.kernel.org/project/linux-scsi/patch/20170212214920.28866-1-
anton@ozlabs.org/#20116693

    Cc: Eric Biederman <ebiederm@xmission.com>
    Signed-off-by: Danielle Ratson <danieller@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 19da1154fc397d372aad351aced0e6182dbaf3ec
Author: Jan Kiszka <jan.kiszka@siemens.com>
Date:   Mon Jun 21 20:08:28 2021 +0200

    soc: ti: pruss: fix referenced node in error message

    [ Upstream commit 8aa35e0bb5eaa42bac415ad0847985daa7b4890c ]

    So far, "(null)" is reported for the node that is missing clocks.

    Signed-off-by: Jan Kiszka <jan.kiszka@siemens.com>
    Acked-by: Suman Anna <s-anna@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/d6e24953-ea89-fd1c-6e16-7a0142118054@siemens.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 620c32a9af98fec55f9f22e2dbeff10824c909e6
Author: Alex Deucher <alexander.deucher@amd.com>
Date:   Wed Nov 10 10:23:25 2021 -0500

    drm/amdgpu/display: set vblank_disable_immediate for DC

    [ Upstream commit 92020e81ddbeac351ea4a19bcf01743f32b9c800 ]

    Disable vblanks immediately to save power.  I think this was
    missed when we merged DC support.

    Bug: https://gitlab.freedesktop.org/drm/amd/-/issues/1781
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2f0fd2f941e88cfb56f1aa5e5ec7bd396576d2f3
Author: Yang Li <yang.lee@linux.alibaba.com>
Date:   Mon Nov 15 16:10:19 2021 +0800

    drm/amd/display: check top_pipe_to_program pointer

    [ Upstream commit a689e8d1f80012f90384ebac9dcfac4201f9f77e ]

    Clang static analysis reports this error

    drivers/gpu/drm/amd/amdgpu/../display/dc/core/dc.c:2870:7: warning:
    Dereference of null pointer [clang-analyzer-core.NullDereference]
                    if
    (top_pipe_to_program->stream_res.tg->funcs->lock_doublebuffer_enable) {
                       ^

    top_pipe_to_program being NULL is caught as an error
    But then it is used to report the error.
```

```
    So add a check before using it.

    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Signed-off-by: Yang Li <yang.lee@linux.alibaba.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8dc363cc9e25c8162ba545057b9464b972c17b48
Author: Anilkumar Kolli <akolli@codeaurora.org>
Date:   Mon Nov 22 13:13:58 2021 +0200

    ath11k: Fix mon status ring rx tlv processing

    [ Upstream commit 09f16f7390f302937409738d6cb6ce99b265f455 ]

    In HE monitor capture, HAL_TLV_STATUS_PPDU_DONE is received
    on processing multiple skb. Do not clear the ppdu_info
    till the HAL_TLV_STATUS_PPDU_DONE is received.

    This fixes below warning and packet drops in monitor mode.
     "Rate marked as an HE rate but data is invalid: MCS: 6, NSS: 0"
     WARNING: at
     PC is at ieee80211_rx_napi+0x624/0x840 [mac80211]

    Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.4.0.1-01693-QCAHKSWPL_SILICONZ-1

    Signed-off-by: Anilkumar Kolli <akolli@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/1637249433-10316-1-git-send-email-akolli@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5354ec4637d648557eed4ee4ddccaa6e71a488f
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Thu Oct 28 16:19:33 2021 +0200

    ARM: imx: rename DEBUG_IMX21_IMX27_UART to DEBUG_IMX27_UART

    [ Upstream commit b0100bce4ff82ec1ccd3c1f3d339fd2df6a81784 ]

    Since commit 4b563a066611 ("ARM: imx: Remove imx21 support"), the config
    DEBUG_IMX21_IMX27_UART is really only debug support for IMX27.

    So, rename this option to DEBUG_IMX27_UART and adjust dependencies in
    Kconfig and rename the definitions to IMX27 as further clean-up.

    This issue was discovered with ./scripts/checkkconfigsymbols.py, which
    reported that DEBUG_IMX21_IMX27_UART depends on the non-existing config
    SOC_IMX21.

    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fd6150e5c0821ea1c6c4ab3e46566e238d80180e
Author: Marek Vasut <marex@denx.de>
Date:   Sat Oct 16 23:05:47 2021 +0200

    soc: imx: gpcv2: Synchronously suspend MIX domains

    [ Upstream commit f756f435f7dd823f2d4bd593ce1bf3168def1308 ]

    In case the following power domain sequence happens, iMX8M Mini always hangs:
      gpumix:on -> gpu:on -> gpu:off -> gpu:on
```

```
    This is likely due to another quirk of the GPC block. This situation can be
    prevented by always synchronously powering off both the domain and MIX domain.
    Make it so. This turns the aforementioned sequence into:
      gpumix:on -> gpu:on -> gpu:off -> gpumix:off -> gpumix:on -> gpu:on

    Signed-off-by: Marek Vasut <marex@denx.de>
    Cc: Frieder Schrempf <frieder.schrempf@kontron.de>
    Cc: Lucas Stach <l.stach@pengutronix.de>
    Cc: NXP Linux Team <linux-imx@nxp.com>
    Cc: Peng Fan <peng.fan@nxp.com>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Acked-by: Lucas Stach <l.stach@pengutronix.de>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be708b4750f2c3d087bd12b406c10bfe5ad4ce3c
Author: Konrad Dybcio <konrad.dybcio@somainline.org>
Date:   Sun Nov 14 02:27:45 2021 +0100

    arm64: dts: qcom: sm8350: Shorten camera-thermal-bottom name

    [ Upstream commit f52dd33943ca5f84ae76890f352f6d9e12512c3f ]

    Thermal zone names should not be longer than 20 names, which is indicated by
    a message at boot. Change "camera-thermal-bottom" to "cam-thermal-bottom" to
    fix it.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@somainline.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211114012755.112226-6-konrad.dybcio@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0795c049a67cb7d094cce83371b45bc46327397c
Author: Konrad Dybcio <konrad.dybcio@somainline.org>
Date:   Sun Nov 14 02:27:44 2021 +0100

    arm64: dts: qcom: sm[68]350: Use interrupts-extended with pdc interrupts

    [ Upstream commit 9e7f7b65c7f04c5cfda97d6bd0d452a49e60f24e ]

    Using interrupts = <&pdc X Y> makes the interrupt framework interpret this as
    the &pdc-nth range of the main interrupt controller (GIC). Fix it.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@somainline.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211114012755.112226-5-konrad.dybcio@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3edfa0b3a9784033f194de739cb71c476f93b913
Author: Dinh Nguyen <dinguyen@kernel.org>
Date:   Tue Oct 12 14:07:06 2021 -0500

    EDAC/synopsys: Use the quirk for version instead of ddr version

    [ Upstream commit bd1d6da17c296bd005bfa656952710d256e77dd3 ]

    Version 2.40a supports DDR_ECC_INTR_SUPPORT for a quirk, so use that
    quirk to determine a call to setup_address_map().

    Signed-off-by: Dinh Nguyen <dinguyen@kernel.org>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Michal Simek <michal.simek@xilinx.com>
    Link: https://lkml.kernel.org/r/20211012190709.1504152-1-dinguyen@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 1bada50b5bfeb87eeeb327ade779587cfcd78ec7
Author: Yang Li <yang.lee@linux.alibaba.com>
Date:   Thu Nov 18 17:48:03 2021 +0800

    ethernet: renesas: Use div64_ul instead of do_div

    [ Upstream commit d9f31aeaa1e5aefa68130878af3c3513d41c1e2d ]

    do_div() does a 64-by-32 division. Here the divisor is an
    unsigned long which on some platforms is 64 bit wide. So use
    div64_ul instead of do_div to avoid a possible truncation.

    Eliminate the following coccicheck warning:
    ./drivers/net/ethernet/renesas/ravb_main.c:2492:1-7: WARNING:
    do_div() does a 64-by-32 division, please consider using div64_ul
    instead.

    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Signed-off-by: Yang Li <yang.lee@linux.alibaba.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Sergey Shtylyov <s.shtylyov@omp.ru>
    Link: https://lore.kernel.org/r/1637228883-100100-1-git-send-email-
yang.lee@linux.alibaba.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6458d123888cc5fe7acca0139599bf9f5f9a7d79
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Nov 17 11:41:13 2021 -0800

    libbpf: Accommodate DWARF/compiler bug with duplicated structs

    [ Upstream commit efdd3eb8015e7447095f02a26eaabd164cd18004 ]

    According to [0], compilers sometimes might produce duplicate DWARF
    definitions for exactly the same struct/union within the same
    compilation unit (CU). We've had similar issues with identical arrays
    and handled them with a similar workaround in 6b6e6b1d09aa ("libbpf:
    Accomodate DWARF/compiler bug with duplicated identical arrays"). Do the
    same for struct/union by ensuring that two structs/unions are exactly
    the same, down to the integer values of field referenced type IDs.

    Solving this more generically (allowing referenced types to be
    equivalent, but using different type IDs, all within a single CU)
    requires a huge complexity increase to handle many-to-many mappings
    between canonidal and candidate type graphs. Before we invest in that,
    let's see if this approach handles all the instances of this issue in
    practice. Thankfully it's pretty rare, it seems.

      [0] https://lore.kernel.org/bpf/YXr2NFlJTAhHdZqq@krava/

    Reported-by: Jiri Olsa <jolsa@kernel.org>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211117194114.347675-1-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffa0e17a1ea85fbc222d991dbf861088dd736e70
Author: Zheyu Ma <zheyuma97@gmail.com>
Date:   Tue May 11 10:00:03 2021 +0100

    media: b2c2: Add missing check in flexcop_pci_isr:

    [ Upstream commit b13203032e679674c7c518f52a7ec0801ca3a829 ]
```

```
    A out-of-bounds bug can be triggered by an interrupt, the reason for
    this bug is the lack of checking of register values.

    In flexcop_pci_isr, the driver reads value from a register and uses it as
    a dma address. Finally, this address will be passed to the count parameter
    of find_next_packet. If this value is larger than the size of dma, the
    index of buffer will be out-of-bounds.

    Fix this by adding a check after reading the value of the register.

    The following KASAN report reveals it:

    BUG: KASAN: slab-out-of-bounds in find_next_packet
    drivers/media/dvb-core/dvb_demux.c:528 [inline]
    BUG: KASAN: slab-out-of-bounds in _dvb_dmx_swfilter
    drivers/media/dvb-core/dvb_demux.c:572 [inline]
    BUG: KASAN: slab-out-of-bounds in dvb_dmx_swfilter+0x3fa/0x420
    drivers/media/dvb-core/dvb_demux.c:603
    Read of size 1 at addr ffff8880608c00a0 by task swapper/2/0

    CPU: 2 PID: 0 Comm: swapper/2 Not tainted 4.19.177-gdba4159c14ef #25
    Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS
    rel-1.12.0-59-gc9ba5276e321-prebuilt.qemu.org 04/01/2014
    Call Trace:
     <IRQ>
     __dump_stack lib/dump_stack.c:77 [inline]
     dump_stack+0xec/0x156 lib/dump_stack.c:118
     print_address_description+0x78/0x290 mm/kasan/report.c:256
     kasan_report_error mm/kasan/report.c:354 [inline]
     kasan_report+0x25b/0x380 mm/kasan/report.c:412
     __asan_report_load1_noabort+0x19/0x20 mm/kasan/report.c:430
     find_next_packet drivers/media/dvb-core/dvb_demux.c:528 [inline]
     _dvb_dmx_swfilter drivers/media/dvb-core/dvb_demux.c:572 [inline]
     dvb_dmx_swfilter+0x3fa/0x420 drivers/media/dvb-core/dvb_demux.c:603
     flexcop_pass_dmx_data+0x2e/0x40 drivers/media/common/b2c2/flexcop.c:167
     flexcop_pci_isr+0x3d1/0x5d0 drivers/media/pci/b2c2/flexcop-pci.c:212
     __handle_irq_event_percpu+0xfb/0x770 kernel/irq/handle.c:149
     handle_irq_event_percpu+0x79/0x150 kernel/irq/handle.c:189
     handle_irq_event+0xac/0x140 kernel/irq/handle.c:206
     handle_fasteoi_irq+0x232/0x5c0 kernel/irq/chip.c:725
     generic_handle_irq_desc include/linux/irqdesc.h:155 [inline]
     handle_irq+0x230/0x3a0 arch/x86/kernel/irq_64.c:87
     do_IRQ+0xa7/0x1e0 arch/x86/kernel/irq.c:247
     common_interrupt+0xf/0xf arch/x86/entry/entry_64.S:670
     </IRQ>
    RIP: 0010:native_safe_halt+0x28/0x30 arch/x86/include/asm/irqflags.h:61
    Code: 00 00 55 be 04 00 00 00 48 c7 c7 00 62 2f 8c 48 89 e5 e8 fb 31
    e8 f8 8b 05 75 4f 8e 03 85 c0 7e 07 0f 00 2d 8a 61 66 00 fb f4 <5d> c3
    90 90 90 90 90 90 0f 1f 44 00 00 55 48 89 e5 41 57 41 56 41
    RSP: 0018:ffff88806b71fcc8 EFLAGS: 00000246 ORIG_RAX: fffffffffffffde
    RAX: 0000000000000000 RBX: ffffffff8bde44c8 RCX: ffffffff88a11285
    RDX: 0000000000000000 RSI: 0000000000000004 RDI: ffffffff8c2f6200
    RBP: ffff88806b71fcc8 R08: fffffbfff185ec40 R09: fffffbfff185ec40
    R10: 0000000000000001 R11: fffffbfff185ec40 R12: 0000000000000002
    R13: ffffffff8be9d6e0 R14: 0000000000000000 R15: 0000000000000000
     arch_safe_halt arch/x86/include/asm/paravirt.h:94 [inline]
     default_idle+0x6f/0x360 arch/x86/kernel/process.c:557
     arch_cpu_idle+0xf/0x20 arch/x86/kernel/process.c:548
     default_idle_call+0x3b/0x60 kernel/sched/idle.c:93
     cpuidle_idle_call kernel/sched/idle.c:153 [inline]
     do_idle+0x2ab/0x3c0 kernel/sched/idle.c:263
     cpu_startup_entry+0xcb/0xe0 kernel/sched/idle.c:369
     start_secondary+0x3b8/0x4e0 arch/x86/kernel/smpboot.c:271
     secondary_startup_64+0xa4/0xb0 arch/x86/kernel/head_64.S:243
```

```
     Allocated by task 1:
      save_stack+0x43/0xd0 mm/kasan/kasan.c:448
      set_track mm/kasan/kasan.c:460 [inline]
      kasan_kmalloc+0xad/0xe0 mm/kasan/kasan.c:553
      kasan_slab_alloc+0x11/0x20 mm/kasan/kasan.c:490
      slab_post_alloc_hook mm/slab.h:445 [inline]
      slab_alloc_node mm/slub.c:2741 [inline]
      slab_alloc mm/slub.c:2749 [inline]
      kmem_cache_alloc+0xeb/0x280 mm/slub.c:2754
      kmem_cache_zalloc include/linux/slab.h:699 [inline]
      __kernfs_new_node+0xe2/0x6f0 fs/kernfs/dir.c:633
      kernfs_new_node+0x9a/0x120 fs/kernfs/dir.c:693
      __kernfs_create_file+0x5f/0x340 fs/kernfs/file.c:992
      sysfs_add_file_mode_ns+0x22a/0x4e0 fs/sysfs/file.c:306
      create_files fs/sysfs/group.c:63 [inline]
      internal_create_group+0x34e/0xc30 fs/sysfs/group.c:147
      sysfs_create_group fs/sysfs/group.c:173 [inline]
      sysfs_create_groups+0x9c/0x140 fs/sysfs/group.c:200
      driver_add_groups+0x3e/0x50 drivers/base/driver.c:129
      bus_add_driver+0x3a5/0x790 drivers/base/bus.c:684
      driver_register+0x1cd/0x410 drivers/base/driver.c:170
      __pci_register_driver+0x197/0x200 drivers/pci/pci-driver.c:1411
      cx88_audio_pci_driver_init+0x23/0x25 drivers/media/pci/cx88/cx88-alsa.c:
      1017
      do_one_initcall+0xe0/0x610 init/main.c:884
      do_initcall_level init/main.c:952 [inline]
      do_initcalls init/main.c:960 [inline]
      do_basic_setup init/main.c:978 [inline]
      kernel_init_freeable+0x4d0/0x592 init/main.c:1145
      kernel_init+0x18/0x190 init/main.c:1062
      ret_from_fork+0x24/0x30 arch/x86/entry/entry_64.S:415

     Freed by task 0:
     (stack is not available)

     The buggy address belongs to the object at ffff8880608c0000
      which belongs to the cache kernfs_node_cache of size 160
     The buggy address is located 0 bytes to the right of
      160-byte region [ffff8880608c0000, ffff8880608c00a0)
     The buggy address belongs to the page:
     page:ffffea0001823000 count:1 mapcount:0 mapping:ffff88806bed1e00
     index:0x0 compound_mapcount: 0
     flags: 0x100000000008100(slab|head)
     raw: 0100000000008100 dead000000000100 dead000000000200 ffff88806bed1e00
     raw: 0000000000000000 0000000000240024 00000001ffffffff 0000000000000000
     page dumped because: kasan: bad access detected

     Memory state around the buggy address:
      ffff8880608bff80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
      ffff8880608c0000: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
     >ffff8880608c0080: 00 00 00 00 fc fc fc fc fc fc fc fc 00 00 00 00
                                    ^
      ffff8880608c0100: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
      ffff8880608c0180: fc fc fc fc fc fc fc fc 00 00 00 00 00 00 00 00
     ==================================================================

     Link: https://lore.kernel.org/linux-media/1620723603-30912-1-git-send-email-
  zheyuma97@gmail.com
     Reported-by: Zheyu Ma <zheyuma97@gmail.com>
     Signed-off-by: Zheyu Ma <zheyuma97@gmail.com>
     Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 97332bf38dd3e7cba499ac4869fa698ba1a8a8a0
  Author: José Expósito <jose.exposito89@gmail.com>
```

```
Date:   Thu Nov 18 08:29:53 2021 +0100

    HID: apple: Do not reset quirks when the Fn key is not found

    [ Upstream commit a5fe7864d8ada170f19cc47d176bf8260ffb4263 ]

    When a keyboard without a function key is detected, instead of removing
    all quirks, remove only the APPLE_HAS_FN quirk.

    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c394bd1bc8537e61593b6b6799e01495c7cf9008
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Thu Nov 18 17:52:08 2021 +0100

    HID: magicmouse: Report battery level over USB

    [ Upstream commit 0b91b4e4dae63cd43871fc2012370b86ee588f91 ]

    When connected over USB, the Apple Magic Mouse 2 and the Apple Magic
    Trackpad 2 register multiple interfaces, one of them is used to report
    the battery level.

    However, unlike when connected over Bluetooth, the battery level is not
    reported automatically and it is required to fetch it manually.

    Fix the battery report descriptor and add a timer to fetch the battery
    level.

    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d4f7ee7716d49eba04a943284530affdaf78f07
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sat Nov 6 14:02:27 2021 +0100

    drm: panel-orientation-quirks: Add quirk for the Lenovo Yoga Book X91F/L

    [ Upstream commit bc30c3b0c8a1904d83d5f0d60fb8650a334b207b ]

    The Lenovo Yoga Book X91F/L uses a panel which has been mounted
    90 degrees rotated. Add a quirk for this.

    Cc: Yauhen Kharuzhy <jekhor@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Acked-by: Simon Ser <contact@emersion.fr>
    Tested-by: Yauhen Kharuzhy <jekhor@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211106130227.11927-1-
hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 415e332497408dfb46608ec3341949bfdebf0f08
Author: Brian Chen <brianchen118@gmail.com>
Date:   Wed Nov 10 21:33:12 2021 +0000

    psi: Fix PSI_MEM_FULL state when tasks are in memstall and doing reclaim

    [ Upstream commit cb0e52b7748737b2cf6481fdd9b920ce7e1ebbdf ]

    We've noticed cases where tasks in a cgroup are stalled on memory but
    there is little memory FULL pressure since tasks stay on the runqueue
    in reclaim.
```

```
    A simple example involves a single threaded program that keeps leaking
    and touching large amounts of memory. It runs in a cgroup with swap
    enabled, memory.high set at 10M and cpu.max ratio set at 5%. Though
    there is significant CPU pressure and memory SOME, there is barely any
    memory FULL since the task enters reclaim and stays on the runqueue.
    However, this memory-bound task is effectively stalled on memory and
    we expect memory FULL to match memory SOME in this scenario.

    The code is confused about memstall && running, thinking there is a
    stalled task and a productive task when there's only one task: a
    reclaimer that's counted as both. To fix this, we redefine the
    condition for PSI_MEM_FULL to check that all running tasks are in an
    active memstall instead of checking that there are no running tasks.

            case PSI_MEM_FULL:
    -               return unlikely(tasks[NR_MEMSTALL] && !tasks[NR_RUNNING]);
    +               return unlikely(tasks[NR_MEMSTALL] &&
    +                       tasks[NR_RUNNING] == tasks[NR_MEMSTALL_RUNNING]);

    This will capture reclaimers. It will also capture tasks that called
    psi_memstall_enter() and are about to sleep, but this should be
    negligible noise.

    Signed-off-by: Brian Chen <brianchen118@gmail.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Acked-by: Johannes Weiner <hannes@cmpxchg.org>
    Link: https://lore.kernel.org/r/20211110213312.310243-1-brianchen118@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ac1fcd2645406099ac14139ab0428d8376c2371
Author: Pavankumar Kondeti <quic_pkondeti@quicinc.com>
Date:   Fri Nov 12 15:54:40 2021 +0530

    usb: gadget: f_fs: Use stream_open() for endpoint files

    [ Upstream commit c76ef96fc00eb398c8fc836b0eb2f82bcc619dc7 ]

    Function fs endpoint file operations are synchronized via an interruptible
    mutex wait. However we see threads that do ep file operations concurrently
    are getting blocked for the mutex lock in __fdget_pos(). This is an
    uninterruptible wait and we see hung task warnings and kernel panic
    if hung_task_panic systcl is enabled if host does not send/receive
    the data for long time.

    The reason for threads getting blocked in __fdget_pos() is due to
    the file position protection introduced by the commit 9c225f2655e3
    ("vfs: atomic f_pos accesses as per POSIX"). Since function fs
    endpoint files does not have the notion of the file position, switch
    to the stream mode. This will bypass the file position mutex and
    threads will be blocked in interruptible state for the function fs
    mutex.

    It should not affects user space as we are only changing the task state
    changes the task state from UNINTERRUPTIBLE to INTERRUPTIBLE while waiting
    for the USB transfers to be finished. However there is a slight change to
    the O_NONBLOCK behavior. Earlier threads that are using O_NONBLOCK are also
    getting blocked inside fdget_pos(). Now they reach to function fs and error
    code is returned. The non blocking behavior is actually honoured now.

    Reviewed-by: John Keeping <john@metanate.com>
    Signed-off-by: Pavankumar Kondeti <quic_pkondeti@quicinc.com>
    Link: https://lore.kernel.org/r/1636712682-1226-1-git-send-email-
quic_pkondeti@quicinc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

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

commit 473c3a3fa4d81ad58da9cbaf18b929c82cb92289
Author: Haimin Zhang <tcs.kernel@gmail.com>
Date:   Sat Nov 13 11:53:20 2021 -0500

    USB: ehci_brcm_hub_control: Improve port index sanitizing

    [ Upstream commit 9933698f6119886c110750e67c10ac66f12b730f ]

    Due to (wIndex & 0xff) - 1 can get an integer greater than 15, this
    can cause array index to be out of bounds since the size of array
    port_status is 15. This change prevents a possible out-of-bounds
    pointer computation by forcing the use of a valid port number.

    Reported-by: TCS Robot <tcs_robot@tencent.com>
    Signed-off-by: Haimin Zhang <tcs.kernel@gmail.com>
    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Link: https://lore.kernel.org/r/20211113165320.GA59686@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fcdc582bd02a0362f365cc891e15e0837f712bf8
Author: Amjad Ouled-Ameur <aouledameur@baylibre.com>
Date:   Fri Nov 12 17:28:26 2021 +0100

    usb: dwc3: meson-g12a: fix shared reset control use

    [ Upstream commit 4ce3b45704d5ef46fb4b28083c8aba6716fabf3b ]

    reset_control_(de)assert() calls are called on a shared reset line when
    reset_control_reset has been used. This is not allowed by the reset
    framework.

    Use reset_control_rearm() call in suspend() and remove() as a way to state
    that the resource is no longer used, hence the shared reset line
    may be triggered again by other devices. Use reset_control_rearm() also in
    case probe fails after reset() has been called.

    reset_control_rearm() keeps use of triggered_count sane in the reset
    framework, use of reset_control_reset() on shared reset line should be
    balanced with reset_control_rearm().

    Signed-off-by: Amjad Ouled-Ameur <aouledameur@baylibre.com>
    Reported-by: Jerome Brunet <jbrunet@baylibre.com>
    Link: https://lore.kernel.org/r/20211112162827.128319-3-aouledameur@baylibre.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15a8d4829c2998e9d99ef709c49967c6cfba0bcb
Author: Baochen Qiang <bqiang@codeaurora.org>
Date:   Tue Oct 26 09:16:05 2021 +0800

    ath11k: Fix crash caused by uninitialized TX ring

    [ Upstream commit 273703ebdb01b6c5f1aaf4b98fb57b177609055c ]

    Commit 31582373a4a8 ("ath11k: Change number of TCL rings to one for
    QCA6390") avoids initializing the other entries of dp->tx_ring cause
    the corresponding TX rings on QCA6390/WCN6855 are not used, but leaves
    those ring masks in ath11k_hw_ring_mask_qca6390.tx unchanged. Normally
    this is OK because we will only get interrupts from the first TX ring
    on these chips and thus only the first entry of dp->tx_ring is involved.

    In case of one MSI vector, all DP rings share the same IRQ. For each
```

```
    interrupt, all rings have to be checked, which means the other entries
    of dp->tx_ring are involved. However since they are not initialized,
    system crashes.

    Fix this issue by simply removing those ring masks.

    crash stack:
    [  102.907438] BUG: kernel NULL pointer dereference, address: 0000000000000028
    [  102.907447] #PF: supervisor read access in kernel mode
    [  102.907451] #PF: error_code(0x0000) - not-present page
    [  102.907453] PGD 1081f0067 P4D 1081f0067 PUD 1081f1067 PMD 0
    [  102.907460] Oops: 0000 [#1] SMP DEBUG_PAGEALLOC NOPTI
    [  102.907465] CPU: 0 PID: 3511 Comm: apt-check Kdump: loaded Tainted: G            E
5.15.0-rc4-wt-ath+ #20
    [  102.907470] Hardware name: AMD Celadon-RN/Celadon-RN, BIOS RCD1005E 10/08/2020
    [  102.907472] RIP: 0010:ath11k_dp_tx_completion_handler+0x201/0x830 [ath11k]
    [  102.907497] Code: 3c 24 4e 8d ac 37 10 04 00 00 4a 8d bc 37 68 04 00 00 48 89 3c 24
48 63 c8 89 83 84 18 00 00 48 c1 e1 05 48 03 8b 78 18 00 00 <8b> 51 08 89 d6 83 e6 07 89 74
24 24 83 fe 03 74 04 85 f6 75 63 41
    [  102.907501] RSP: 0000:ffff9b7340003e08 EFLAGS: 00010202
    [  102.907505] RAX: 0000000000000001 RBX: ffff8e21530c0100 RCX: 0000000000000020
    [  102.907508] RDX: 0000000000000000 RSI: 00000000fffffe00 RDI: ffff8e21530c1938
    [  102.907511] RBP: ffff8e21530c0000 R08: 0000000000000001 R09: 0000000000000000
    [  102.907513] R10: ffff8e2145534c10 R11: 0000000000000001 R12: ffff8e21530c2938
    [  102.907515] R13: ffff8e21530c18e0 R14: 0000000000000100 R15: ffff8e21530c2978
    [  102.907518] FS:  00007f5d4297e740(0000) GS:ffff8e243d600000(0000)
knlGS:0000000000000000
    [  102.907521] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [  102.907524] CR2: 0000000000000028 CR3: 00000001034ea000 CR4: 0000000000350ef0
    [  102.907527] Call Trace:
    [  102.907531]  <IRQ>
    [  102.907537]  ath11k_dp_service_srng+0x5c/0x2f0 [ath11k]
    [  102.907556]  ath11k_pci_ext_grp_napi_poll+0x21/0x70 [ath11k_pci]
    [  102.907562]  __napi_poll+0x2c/0x160
    [  102.907570]  net_rx_action+0x251/0x310
    [  102.907576]  __do_softirq+0x107/0x2fc
    [  102.907585]  irq_exit_rcu+0x74/0x90
    [  102.907593]  common_interrupt+0x83/0xa0
    [  102.907600]  </IRQ>
    [  102.907601]  asm_common_interrupt+0x1e/0x40

    Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1

    Signed-off-by: Baochen Qiang <bqiang@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211026011605.58615-1-quic_bqiang@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b202a59fc5c92822477a9ea812d8f4ae629d7ffc
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Thu Nov 11 23:04:09 2021 +0000

    media: atomisp: handle errors at sh_css_create_isp_params()

    [ Upstream commit 58043dbf6d1ae9deab4f5aa1e039c70112017682 ]

    The succ var tracks memory allocation erros on this function.

    Fix it, in order to stop this W=1 Werror in clang:

    drivers/staging/media/atomisp/pci/sh_css_params.c:2430:7: error: variable 'succ' set but
not used [-Werror,-Wunused-but-set-variable]
            bool succ = true;
                 ^
```

```
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22f78512b4f1bf920a689eeeb088317c3818efb1
Author: Tiezhu Yang <yangtiezhu@loongson.cn>
Date:   Fri Nov 5 09:30:00 2021 +0800

    bpf: Change value of MAX_TAIL_CALL_CNT from 32 to 33

    [ Upstream commit ebf7f6f0a6cdcc17a3da52b81e4b3a98c4005028 ]

    In the current code, the actual max tail call count is 33 which is greater
    than MAX_TAIL_CALL_CNT (defined as 32). The actual limit is not consistent
    with the meaning of MAX_TAIL_CALL_CNT and thus confusing at first glance.
    We can see the historical evolution from commit 04fd61ab36ec ("bpf: allow
    bpf programs to tail-call other bpf programs") and commit f9dabe016b63
    ("bpf: Undo off-by-one in interpreter tail call count limit"). In order
    to avoid changing existing behavior, the actual limit is 33 now, this is
    reasonable.

    After commit 874be05f525e ("bpf, tests: Add tail call test suite"), we can
    see there exists failed testcase.

    On all archs when CONFIG_BPF_JIT_ALWAYS_ON is not set:
     # echo 0 > /proc/sys/net/core/bpf_jit_enable
     # modprobe test_bpf
     # dmesg | grep -w FAIL
     Tail call error path, max count reached jited:0 ret 34 != 33 FAIL

    On some archs:
     # echo 1 > /proc/sys/net/core/bpf_jit_enable
     # modprobe test_bpf
     # dmesg | grep -w FAIL
     Tail call error path, max count reached jited:1 ret 34 != 33 FAIL

    Although the above failed testcase has been fixed in commit 18935a72eb25
    ("bpf/tests: Fix error in tail call limit tests"), it would still be good
    to change the value of MAX_TAIL_CALL_CNT from 32 to 33 to make the code
    more readable.

    The 32-bit x86 JIT was using a limit of 32, just fix the wrong comments and
    limit to 33 tail calls as the constant MAX_TAIL_CALL_CNT updated. For the
    mips64 JIT, use "ori" instead of "addiu" as suggested by Johan Almbladh.
    For the riscv JIT, use RV_REG_TCC directly to save one register move as
    suggested by Björn Töpel. For the other implementations, no function changes,
    it does not change the current limit 33, the new value of MAX_TAIL_CALL_CNT
    can reflect the actual max tail call count, the related tail call testcases
    in test_bpf module and selftests can work well for the interpreter and the
    JIT.

    Here are the test results on x86_64:

     # uname -m
     x86_64
     # echo 0 > /proc/sys/net/core/bpf_jit_enable
     # modprobe test_bpf test_suite=test_tail_calls
     # dmesg | tail -1
     test_bpf: test_tail_calls: Summary: 8 PASSED, 0 FAILED, [0/8 JIT'ed]
     # rmmod test_bpf
     # echo 1 > /proc/sys/net/core/bpf_jit_enable
     # modprobe test_bpf test_suite=test_tail_calls
     # dmesg | tail -1
     test_bpf: test_tail_calls: Summary: 8 PASSED, 0 FAILED, [8/8 JIT'ed]
     # rmmod test_bpf
```

```
      # ./test_progs -t tailcalls
      #142 tailcalls:OK
      Summary: 1/11 PASSED, 0 SKIPPED, 0 FAILED

    Signed-off-by: Tiezhu Yang <yangtiezhu@loongson.cn>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Tested-by: Johan Almbladh <johan.almbladh@anyfinetworks.com>
    Tested-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Acked-by: Björn Töpel <bjorn@kernel.org>
    Acked-by: Johan Almbladh <johan.almbladh@anyfinetworks.com>
    Acked-by: Ilya Leoshkevich <iii@linux.ibm.com>
    Link: https://lore.kernel.org/bpf/1636075800-3264-1-git-send-email-
yangtiezhu@loongson.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e78b91a5867240a8f190e4d28086f712a7cf2c79
Author: Linus Lüssing <linus.luessing@c0d3.blue>
Date:   Sun Oct 31 22:30:12 2021 +0100

    batman-adv: allow netlink usage in unprivileged containers

    [ Upstream commit 9057d6c23e7388ee9d037fccc9a7bc8557ce277b ]

    Currently, creating a batman-adv interface in an unprivileged LXD
    container and attaching secondary interfaces to it with "ip" or "batctl"
    works fine. However all batctl debug and configuration commands
    fail:

      root@container:~# batctl originators
      Error received: Operation not permitted
      root@container:~# batctl orig_interval
      1000
      root@container:~# batctl orig_interval 2000
      root@container:~# batctl orig_interval
      1000

    To fix this change the generic netlink permissions from GENL_ADMIN_PERM
    to GENL_UNS_ADMIN_PERM. This way a batman-adv interface is fully
    maintainable as root from within a user namespace, from an unprivileged
    container.

    All except one batman-adv netlink setting are per interface and do not
    leak information or change settings from the host system and are
    therefore save to retrieve or modify as root from within an unprivileged
    container.

    "batctl routing_algo" / BATADV_CMD_GET_ROUTING_ALGOS is the only
    exception: It provides the batman-adv kernel module wide default routing
    algorithm. However it is read-only from netlink and an unprivileged
    container is still not allowed to modify
    /sys/module/batman_adv/parameters/routing_algo. Instead it is advised to
    use the newly introduced "batctl if create routing_algo RA_NAME" /
    IFLA_BATADV_ALGO_NAME to set the routing algorithm on interface
    creation, which already works fine in an unprivileged container.

    Cc: Tycho Andersen <tycho@tycho.pizza>
    Signed-off-by: Linus Lüssing <linus.luessing@c0d3.blue>
    Signed-off-by: Sven Eckelmann <sven@narfation.org>
    Signed-off-by: Simon Wunderlich <sw@simonwunderlich.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fc7cbcaa53af21355a40245fc183940e9d6c1a3b
Author: Wen Gong <wgong@codeaurora.org>
Date:   Wed Oct 13 03:37:04 2021 -0400
```

```
    ath11k: enable IEEE80211_VHT_EXT_NSS_BW_CAPABLE if NSS ratio enabled

    [ Upstream commit 78406044bdd0cc8987bc082b76867c63ab1c6af8 ]

    When NSS ratio enabled reported by firmware, SUPPORTS_VHT_EXT_NSS_BW
    is set in ath11k, meanwhile IEEE80211_VHT_EXT_NSS_BW_CAPABLE also
    need to be set, otherwise it is invalid because spec in IEEE Std
    802.11™-2020 as below.

    Table 9-273-Supported VHT-MCS and NSS Set subfields, it has subfield
    VHT Extended NSS BW Capable, its definition is:
    Indicates whether the STA is capable of interpreting the Extended NSS
    BW Support subfield of the VHT Capabilities Information field.

    dmesg have a message without this patch:

    ieee80211 phy0: copying sband (band 1) due to VHT EXT NSS BW flag

    It means mac80211 will set IEEE80211_VHT_EXT_NSS_BW_CAPABLE if ath11k not
    set it in ieee80211_register_hw(). So it is better to set it in ath11k.

    Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1

    Signed-off-by: Wen Gong <wgong@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211013073704.15888-1-wgong@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b48897f0754d56f812026e7906f8c2df7c6f11aa
Author: Wan Jiabing <wanjiabing@vivo.com>
Date:   Sun Oct 17 21:45:03 2021 -0400

    ARM: shmobile: rcar-gen2: Add missing of_node_put()

    [ Upstream commit 85744f2d938c5f3cfc44cb6533c157469634da93 ]

    Fix following coccicheck warning:
    ./arch/arm/mach-shmobile/regulator-quirk-rcar-gen2.c:156:1-33: Function
    for_each_matching_node_and_match should have of_node_put() before break
    and goto.

    Early exits from for_each_matching_node_and_match() should decrement the
    node reference counter.

    Signed-off-by: Wan Jiabing <wanjiabing@vivo.com>
    Link: https://lore.kernel.org/r/20211018014503.7598-1-wanjiabing@vivo.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 399a3994858e501ee0be7dd780ca364e8af04913
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Wed Nov 10 12:59:11 2021 +0000

    media: atomisp: check before deference asd variable

    [ Upstream commit 71665d816214124d6bc4eb80314ac8f84ecacd78 ]

    The asd->isp was referenced before checking if asd is not
    NULL.

    This fixes this warning:

            ../drivers/staging/media/atomisp/pci/atomisp_cmd.c:5548 atomisp_set_fmt_to_snr()
warn: variable dereferenced before check 'asd' (see line 5540)
```

```
    While here, avoid getting the pipe pointer twice.

    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57bee607744275688ce22c2cd80adcceb8de22e2
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Nov 7 17:15:48 2021 +0000

    media: atomisp-ov2680: Fix ov2680_set_fmt() clobbering the exposure

    [ Upstream commit 4492289c31364d28c2680b43b18883385a5d216c ]

    Now that we restore the default or last user set exposure setting on
    power_up() there is no need for the registers written by ov2680_set_fmt()
    to write to the exposure register.

    Not doing so fixes the exposure always being reset to the value from
    the res->regs array after a set_fmt().

    Link: https://lore.kernel.org/linux-media/20211107171549.267583-11-hdegoede@redhat.com
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4177e32307543bca04dd1a1c1d89569e2f8b3d83
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Fri Oct 29 08:09:39 2021 +0100

    media: atomisp: set per-device's default mode

    [ Upstream commit 2c45e343c581091835c9047ed5298518aa133163 ]

    The atomisp driver originally used the s_parm command to
    initialize the run_mode type to the driver. So, before start
    setting up the streaming, s_parm should be called.

    So, even having 5 "normal" video devices, one meant to be used
    for each type, the run_mode was actually selected when
    s_parm is called.

    Without setting the run mode, applications that don't call
    VIDIOC_SET_PARM with a custom atomisp parameters won't work, as
    the pipeline won't be set:

            atomisp-isp2 0000:00:03.0: can't create streams
            atomisp-isp2 0000:00:03.0: __get_frame_info 1600x1200 (padded to 0) returned -22

    However, commit 8a7c5594c020 ("media: v4l2-ioctl: clear fields in s_parm")
    broke support for it, with a good reason, as drivers shoudn't be
    extending the API for their own purposes.

    So, as an step to allow generic apps to use this driver, put
    the device's run_mode in preview after open.

    After this patch, using v4l2grab starts to work on preview
    mode (/dev/video2):

            $ v4l2grab -f YUYV -x 1600 -y 1200 -d /dev/video2 -n 1 -u
            $ feh out000.pnm

    So, let's just setup the default run_mode that each video devnode
    should assume, setting it at open() time.

    Reported-by: Tsuchiya Yuto <kitakar@gmail.com>
```

```
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3696bf77800bb946b8824281e33a9cf4c52290dd
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Thu Nov 4 11:45:27 2021 +0000

    media: atomisp: fix try_fmt logic

    [ Upstream commit c9e9094c4e42124af909b2f5f6ded0498e0854ac ]

    The internal try_fmt logic is not meant to provide everything
    that the V4L2 API should provide. Also, it doesn't decrement
    the pads that are used only internally by the driver, but aren't
    part of the device's output.

    Fix it.

    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b98913f3d035f639eda2e213e308fd5567c00d2
Author: Ben Skeggs <bskeggs@redhat.com>
Date:   Wed Feb 24 19:29:52 2021 +1000

    drm/nouveau/pmu/gm200-: avoid touching PMU outside of DEVINIT/PREOS/ACR

    [ Upstream commit 1d2271d2fb85e54bfc9630a6c30ac0feb9ffb983 ]

    There have been reports of the WFI timing out on some boards, and a
    patch was proposed to just remove it.  This stuff is rather fragile,
    and I believe the WFI might be needed with our FW prior to GM200.

    However, we probably should not be touching PMU during init on GPUs
    where we depend on NVIDIA FW, outside of limited circumstances, so
    this should be a somewhat safer change that achieves the desired
    result.

    Reported-by: Diego Viola <diego.viola@gmail.com>
    Signed-off-by: Ben Skeggs <bskeggs@redhat.com>
    Reviewed-by: Karol Herbst <kherbst@redhat.com>
    Signed-off-by: Karol Herbst <kherbst@redhat.com>
    Link: https://gitlab.freedesktop.org/drm/nouveau/-/merge_requests/10
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f21c22fbcc54b0aea05d7e5795e1c59e30ad19e6
Author: Neil Armstrong <narmstrong@baylibre.com>
Date:   Fri Oct 29 15:59:47 2021 +0200

    drm/bridge: dw-hdmi: handle ELD when DRM_BRIDGE_ATTACH_NO_CONNECTOR

    [ Upstream commit 3f2532d65a571ca02258b547b5b68ab2e9406fdb ]

    The current ELD handling takes the internal connector ELD buffer and
    shares it to the I2S and AHB sub-driver.

    But with DRM_BRIDGE_ATTACH_NO_CONNECTOR, the connector is created
    elsewhere (or not), and an eventual connector is known only
    if the bridge chain up to a connector is enabled.

    The current dw-hdmi code gets the current connector from
    atomic_enable() so use the already stored connector pointer and
    replace the buffer pointer with a callback returning the current
    connector ELD buffer.
```

```
    Since a connector is not always available, either pass an empty
    ELD to the alsa HDMI driver or don't call snd_pcm_hw_constraint_eld()
    in AHB driver.

    Reported-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    [narmstrong: fixed typo in commit log]
    Acked-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211029135947.3022875-1-
narmstrong@baylibre.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c91e7952bce2d5293503a0466777b44682017ebc
Author: Zekun Shen <bruceshenzk@gmail.com>
Date:   Thu Oct 28 18:37:49 2021 -0400

    ar5523: Fix null-ptr-deref with unexpected WDCMSG_TARGET_START reply

    [ Upstream commit ae80b6033834342601e99f74f6a62ff5092b1cee ]

    Unexpected WDCMSG_TARGET_START replay can lead to null-ptr-deref
    when ar->tx_cmd->odata is NULL. The patch adds a null check to
    prevent such case.

    KASAN: null-ptr-deref in range [0x0000000000000000-0x0000000000000007]
     ar5523_cmd+0x46a/0x581 [ar5523]
     ar5523_probe.cold+0x1b7/0x18da [ar5523]
     ? ar5523_cmd_rx_cb+0x7a0/0x7a0 [ar5523]
     ? __pm_runtime_set_status+0x54a/0x8f0
     ? _raw_spin_trylock_bh+0x120/0x120
     ? pm_runtime_barrier+0x220/0x220
     ? __pm_runtime_resume+0xb1/0xf0
     usb_probe_interface+0x25b/0x710
     really_probe+0x209/0x5d0
     driver_probe_device+0xc6/0x1b0
     device_driver_attach+0xe2/0x120

    I found the bug using a custome USBFuzz port. It's a research work
    to fuzz USB stack/drivers. I modified it to fuzz ath9k driver only,
    providing hand-crafted usb descriptors to QEMU.

    After fixing the code (fourth byte in usb packet) to WDCMSG_TARGET_START,
    I got the null-ptr-deref bug. I believe the bug is triggerable whenever
    cmd->odata is NULL. After patching, I tested with the same input and no
    longer see the KASAN report.

    This was NOT tested on a real device.

    Signed-off-by: Zekun Shen <bruceshenzk@gmail.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/YXsmPQ3awHFLuAj2@10-18-43-117.dynapool.wireless.nyu.edu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 11df5857da6f99604cf213a1b3a10eaeca40b341
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Sun Nov 7 08:55:21 2021 -0800

    selftests/bpf: Fix bpf_object leak in skb_ctx selftest

    [ Upstream commit 8c7a95520184b6677ca6075e12df9c208d57d088 ]

    skb_ctx selftest didn't close bpf_object implicitly allocated by
    bpf_prog_test_load() helper. Fix the problem by explicitly calling
    bpf_object__close() at the end of the test.
```

```
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Reviewed-by: Hengqi Chen <hengqi.chen@gmail.com>
    Link: https://lore.kernel.org/bpf/20211107165521.9240-10-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1fc68e1302bfb6b4bfd36fb0d9cbe30fecd5559b
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Sun Nov 7 08:55:20 2021 -0800

    selftests/bpf: Destroy XDP link correctly

    [ Upstream commit f91231eeeed752119f49eb6620cae44ec745a007 ]

    bpf_link__detach() was confused with bpf_link__destroy() and leaves
    leaked FD in the process. Fix the problem.

    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Reviewed-by: Hengqi Chen <hengqi.chen@gmail.com>
    Link: https://lore.kernel.org/bpf/20211107165521.9240-9-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5a6da5536544fee374ae1f04871ea3724ecd1d1
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Sun Nov 7 08:55:15 2021 -0800

    selftests/bpf: Fix memory leaks in btf_type_c_dump() helper

    [ Upstream commit 8ba285874913da21ca39a46376e9cc5ce0f45f94 ]

    Free up memory and resources used by temporary allocated memstream and
    btf_dump instance.

    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Reviewed-by: Hengqi Chen <hengqi.chen@gmail.com>
    Link: https://lore.kernel.org/bpf/20211107165521.9240-4-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 87ab0d4c44740e6dc8df1d3f82c0e6b423651f08
Author: Qiang Yu <yuq825@gmail.com>
Date:   Sun Oct 31 12:16:04 2021 +0800

    drm/lima: fix warning when CONFIG_DEBUG_SG=y & CONFIG_DMA_API_DEBUG=y

    [ Upstream commit 89636a06fa2ee7826a19c39c19a9bc99ab9340a9 ]

    Otherwise get following warning:

    DMA-API: lima 1c40000.gpu: mapping sg segment longer than device claims to support
[len=4149248] [max=65536]

    See: https://gitlab.freedesktop.org/mesa/mesa/-/issues/5496

    Reviewed-by: Vasily Khoruzhick <anarsoul@gmail.com>
    Reported-by: Roman Stratiienko <r.stratiienko@gmail.com>
    Signed-off-by: Qiang Yu <yuq825@gmail.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211031041604.187216-1-
yuq825@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37e64c55232b1d2cd46647d5c781351416abe626
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Nov 3 10:32:10 2021 -0700
```

```
    libbpf: Improve sanity checking during BTF fix up

    [ Upstream commit 88918dc12dc357a06d8d722a684617b1c87a4654 ]

    If BTF is corrupted DATASEC's variable type ID might be incorrect.
    Prevent this easy to detect situation with extra NULL check.
    Reported by oss-fuzz project.

    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Link: https://lore.kernel.org/bpf/20211103173213.1376990-3-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 614e2ce4aea3fd525bb29874599badd2bee90dca
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Nov 3 10:32:09 2021 -0700

    libbpf: Detect corrupted ELF symbols section

    [ Upstream commit 833907876be55205d0ec153dcd819c014404ee16 ]

    Prevent divide-by-zero if ELF is corrupted and has zero sh_entsize.
    Reported by oss-fuzz project.

    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Link: https://lore.kernel.org/bpf/20211103173213.1376990-2-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df1fe4c614f4747b4b34958c37d0ab10c4b79204
Author: Alexander Aring <aahringo@redhat.com>
Date:   Tue Nov 2 15:17:24 2021 -0400

    fs: dlm: filter user dlm messages for kernel locks

    [ Upstream commit 6c2e3bf68f3e5e5a647aa52be246d5f552d7496d ]

    This patch fixes the following crash by receiving a invalid message:

    [  160.672220] ================================================================
    [  160.676206] BUG: KASAN: user-memory-access in dlm_user_add_ast+0xc3/0x370
    [  160.679659] Read of size 8 at addr 00000000deadbeef by task kworker/u32:13/319
    [  160.681447]
    [  160.681824] CPU: 10 PID: 319 Comm: kworker/u32:13 Not tainted 5.14.0-rc2+ #399
    [  160.683472] Hardware name: Red Hat KVM/RHEL-AV, BIOS 1.14.0-
1.module+el8.6.0+12648+6ede71a5 04/01/2014
    [  160.685574] Workqueue: dlm_recv process_recv_sockets
    [  160.686721] Call Trace:
    [  160.687310]  dump_stack_lvl+0x56/0x6f
    [  160.688169]  ? dlm_user_add_ast+0xc3/0x370
    [  160.689116]  kasan_report.cold.14+0x116/0x11b
    [  160.690138]  ? dlm_user_add_ast+0xc3/0x370
    [  160.690832]  dlm_user_add_ast+0xc3/0x370
    [  160.691502]  _receive_unlock_reply+0x103/0x170
    [  160.692241]  _receive_message+0x11df/0x1ec0
    [  160.692926]  ? rcu_read_lock_sched_held+0xa1/0xd0
    [  160.693700]  ? rcu_read_lock_bh_held+0xb0/0xb0
    [  160.694427]  ? lock_acquire+0x175/0x400
    [  160.695058]  ? do_purge.isra.51+0x200/0x200
    [  160.695744]  ? lock_acquired+0x360/0x5d0
    [  160.696400]  ? lock_contended+0x6a0/0x6a0
    [  160.697055]  ? lock_release+0x21d/0x5e0
```

```
[  160.697686]  ? lock_is_held_type+0xe0/0x110
[  160.698352]  ? lock_is_held_type+0xe0/0x110
[  160.699026]  ? ___might_sleep+0x1cc/0x1e0
[  160.699698]  ? dlm_wait_requestqueue+0x94/0x140
[  160.700451]  ? dlm_process_requestqueue+0x240/0x240
[  160.701249]  ? down_write_killable+0x2b0/0x2b0
[  160.701988]  ? do_raw_spin_unlock+0xa2/0x130
[  160.702690]  dlm_receive_buffer+0x1a5/0x210
[  160.703385]  dlm_process_incoming_buffer+0x726/0x9f0
[  160.704210]  receive_from_sock+0x1c0/0x3b0
[  160.704886]  ? dlm_tcp_shutdown+0x30/0x30
[  160.705561]  ? lock_acquire+0x175/0x400
[  160.706197]  ? rcu_read_lock_sched_held+0xa1/0xd0
[  160.706941]  ? rcu_read_lock_bh_held+0xb0/0xb0
[  160.707681]  process_recv_sockets+0x32/0x40
[  160.708366]  process_one_work+0x55e/0xad0
[  160.709045]  ? pwq_dec_nr_in_flight+0x110/0x110
[  160.709820]  worker_thread+0x65/0x5e0
[  160.710423]  ? process_one_work+0xad0/0xad0
[  160.711087]  kthread+0x1ed/0x220
[  160.711628]  ? set_kthread_struct+0x80/0x80
[  160.712314]  ret_from_fork+0x22/0x30
```

The issue is that we received a DLM message for a user lock but the destination lock is a kernel lock. Note that the address which is trying to derefence is 00000000deadbeef, which is in a kernel lock lkb->lkb_astparam, this field should never be derefenced by the DLM kernel stack. In case of a user lock lkb->lkb_astparam is lkb->lkb_ua (memory is shared by a union field). The struct lkb_ua will be handled by the DLM kernel stack but on a kernel lock it will contain invalid data and ends in most likely crashing the kernel.

It can be reproduced with two cluster nodes.

node 2:

```
dlm_tool join test
echo "862 fooobaar 1 2 1" > /sys/kernel/debug/dlm/test_locks
echo "862 3 1" > /sys/kernel/debug/dlm/test_waiters
```

node 1:

```
dlm_tool join test
```

python:

```
foo = DLM(h_cmd=3, o_nextcmd=1, h_nodeid=1, h_lockspace=0x77222027, \
          m_type=7, m_flags=0x1, m_remid=0x862, m_result=0xFFFEFFFE)
newFile = open("/sys/kernel/debug/dlm/comms/2/rawmsg", "wb")
newFile.write(bytes(foo))
```

Signed-off-by: Alexander Aring <aahringo@redhat.com>
Signed-off-by: David Teigland <teigland@redhat.com>
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f7995c4a7e3dd10d1efb32fea2d22aa2a7303c7
Author: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
Date: Thu Oct 28 12:24:03 2021 -0400

drm/sched: Avoid lockdep spalt on killing a processes

[Upstream commit 542cff7893a37445f98ece26aeb3c9c1055e9ea4]

Probelm:
Singlaning one sched fence from within another's sched fence singal callback generates lockdep splat because the both have same lockdep class of their fence->lock

```
    Fix:
    Fix bellow stack by rescheduling to irq work of
    signaling and killing of jobs that left when entity is killed.

    [11176.741181]  dump_stack+0x10/0x12
    [11176.741186] __lock_acquire.cold+0x208/0x2df
    [11176.741197]  lock_acquire+0xc6/0x2d0
    [11176.741204]  ? dma_fence_signal+0x28/0x80
    [11176.741212] _raw_spin_lock_irqsave+0x4d/0x70
    [11176.741219]  ? dma_fence_signal+0x28/0x80
    [11176.741225]  dma_fence_signal+0x28/0x80
    [11176.741230] drm_sched_fence_finished+0x12/0x20 [gpu_sched]
    [11176.741240] drm_sched_entity_kill_jobs_cb+0x1c/0x50 [gpu_sched]
    [11176.741248] dma_fence_signal_timestamp_locked+0xac/0x1a0
    [11176.741254]  dma_fence_signal+0x3b/0x80
    [11176.741260] drm_sched_fence_finished+0x12/0x20 [gpu_sched]
    [11176.741268] drm_sched_job_done.isra.0+0x7f/0x1a0 [gpu_sched]
    [11176.741277] drm_sched_job_done_cb+0x12/0x20 [gpu_sched]
    [11176.741284] dma_fence_signal_timestamp_locked+0xac/0x1a0
    [11176.741290]  dma_fence_signal+0x3b/0x80
    [11176.741296] amdgpu_fence_process+0xd1/0x140 [amdgpu]
    [11176.741504] sdma_v4_0_process_trap_irq+0x8c/0xb0 [amdgpu]
    [11176.741731]  amdgpu_irq_dispatch+0xce/0x250 [amdgpu]
    [11176.741954]  amdgpu_ih_process+0x81/0x100 [amdgpu]
    [11176.742174]  amdgpu_irq_handler+0x26/0xa0 [amdgpu]
    [11176.742393] __handle_irq_event_percpu+0x4f/0x2c0
    [11176.742402] handle_irq_event_percpu+0x33/0x80
    [11176.742408]  handle_irq_event+0x39/0x60
    [11176.742414]  handle_edge_irq+0x93/0x1d0
    [11176.742419]  __common_interrupt+0x50/0xe0
    [11176.742426]  common_interrupt+0x80/0x90

    Signed-off-by: Andrey Grodzovsky <andrey.grodzovsky@amd.com>
    Suggested-by: Daniel Vetter  <daniel.vetter@ffwll.ch>
    Suggested-by: Christian König <christian.koenig@amd.com>
    Tested-by: Christian König <christian.koenig@amd.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Link: https://www.spinics.net/lists/dri-devel/msg321250.html
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0406b8bd5652d92ea1fa354fab2f2b537c1880b4
Author: Archie Pusaka <apusaka@chromium.org>
Date:   Thu Oct 28 19:17:25 2021 +0800

    Bluetooth: Fix removing adv when processing cmd complete

    [ Upstream commit 2128939fe2e771645dd88e1938c27fdf96bd1cd0 ]

    If we remove one instance of adv using Set Extended Adv Enable, there
    is a possibility of issue occurs when processing the Command Complete
    event. Especially, the adv_info might not be found since we already
    remove it in hci_req_clear_adv_instance() -> hci_remove_adv_instance().
    If that's the case, we will mistakenly proceed to remove all adv
    instances instead of just one single instance.

    This patch fixes the issue by checking the content of the HCI command
    instead of checking whether the adv_info is found.

    Signed-off-by: Archie Pusaka <apusaka@chromium.org>
    Reviewed-by: Sonny Sasaka <sonnysasaka@chromium.org>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 27053684f93ff909af388e6c24a630f58c6217b5
Author: Brian Norris <briannorris@chromium.org>
```

```
Date:   Thu Sep 23 17:33:55 2021 -0700

    drm/panel: Delete panel on mipi_dsi_attach() failure

    [ Upstream commit 9bf7123bb07f98dc76acb5daa91248e6f95713cb ]

    Many DSI panel drivers fail to clean up their panel references on
    mipi_dsi_attach() failure, so we're leaving a dangling drm_panel
    reference to freed memory. Clean that up on failure.

    Noticed by inspection, after seeing similar problems on other drivers.
    Therefore, I'm not marking Fixes/stable.

    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210923173336.3.If9e74fa9b1d6eaa9e0e5b95b2b95
7b992740251c@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03b036ba1017713fb23d2276cc54232a77f15ca3
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Wed Oct 13 16:55:01 2021 +0800

    Bluetooth: Fix memory leak of hci device

    [ Upstream commit 75d9b8559ac36e059238ee4f8e33cd86086586ba ]

    Fault injection test reported memory leak of hci device as follows:

    unreferenced object 0xffff88800b858000 (size 8192):
      comm "kworker/0:2", pid 167, jiffies 4294955747 (age 557.148s)
      hex dump (first 32 bytes):
        00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  ................
        00 00 00 00 00 00 00 00 00 00 00 00 ad 4e ad de  .............N..
      backtrace:
        [<0000000070eb1059>] kmem_cache_alloc_trace mm/slub.c:3208
        [<00000000015eb521>] hci_alloc_dev_priv include/linux/slab.h:591
        [<00000000dcfc1e21>] bpa10x_probe include/net/bluetooth/hci_core.h:1240
        [<000000005d3028c7>] usb_probe_interface drivers/usb/core/driver.c:397
        [<00000000cbac9243>] really_probe drivers/base/dd.c:517
        [<0000000024cab3f0>] __driver_probe_device drivers/base/dd.c:751
        [<00000000202135cb>] driver_probe_device drivers/base/dd.c:782
        [<000000000761f2bc>] __device_attach_driver drivers/base/dd.c:899
        [<00000000f7d63134>] bus_for_each_drv drivers/base/bus.c:427
        [<00000000c9551f0b>] __device_attach drivers/base/dd.c:971
        [<000000007f79bd16>] bus_probe_device drivers/base/bus.c:487
        [<000000007bb8b95a>] device_add drivers/base/core.c:3364
        [<000000009564d9ea>] usb_set_configuration drivers/usb/core/message.c:2171
        [<00000000e4657087>] usb_generic_driver_probe drivers/usb/core/generic.c:239
        [<0000000071ede518>] usb_probe_device drivers/usb/core/driver.c:294
        [<00000000cbac9243>] really_probe drivers/base/dd.c:517

    hci_alloc_dev() do not init the device's flag. And hci_free_dev()
    using put_device() to free the memory allocated for this device,
    but it calls just put_device(dev) only in case of HCI_UNREGISTER
    flag is set, So any error handing before hci_register_dev() success
    will cause memory leak.

    To avoid this behaviour we can using kfree() to release dev before
    hci_register_dev() success.

    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 139fee4f46bf1ea83d8bd3b1bc65e044d807327b
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Wed Oct 13 16:55:46 2021 +0800

    Bluetooth: Fix debugfs entry leak in hci_register_dev()

    [ Upstream commit 5a4bb6a8e981d3d0d492aa38412ee80b21033177 ]

    Fault injection test report debugfs entry leak as follows:

    debugfs: Directory 'hci0' with parent 'bluetooth' already present!

    When register_pm_notifier() failed in hci_register_dev(), the debugfs
    create by debugfs_create_dir() do not removed in the error handing path.

    Add the remove debugfs code to fix it.

    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 324d8565709bae00be1060983460b6d0a077c7e5
Author: Nguyen Dinh Phi <phind.uet@gmail.com>
Date:   Fri Oct 8 03:04:24 2021 +0800

    Bluetooth: hci_sock: purge socket queues in the destruct() callback

    [ Upstream commit 709fca500067524381e28a5f481882930eebac88 ]

    The receive path may take the socket right before hci_sock_release(),
    but it may enqueue the packets to the socket queues after the call to
    skb_queue_purge(), therefore the socket can be destroyed without clear
    its queues completely.

    Moving these skb_queue_purge() to the hci_sock_destruct() will fix this
    issue, because nothing is referencing the socket at this point.

    Signed-off-by: Nguyen Dinh Phi <phind.uet@gmail.com>
    Reported-by: syzbot+4c4ffd1e1094dae61035@syzkaller.appspotmail.com
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3019d99bc4de3a7a90ab5fdd664530019996706
Author: Merlijn Wajer <merlijn@wizzup.org>
Date:   Sun Dec 12 23:40:07 2021 +0100

    leds: lp55xx: initialise output direction from dts

    [ Upstream commit 9e87a8da747bf72365abb79e6f64fcca955b4f56 ]

    Commit a5d3d1adc95f ("leds: lp55xx: Initialize enable GPIO direction to
    output") attempts to fix this, but the fix did not work since at least
    for the Nokia N900 the value needs to be set to HIGH, per the device
    tree. So rather than hardcoding the value to a potentially invalid value
    for some devices, let's set direction in lp55xx_init_device.

    Fixes: a5d3d1adc95f ("leds: lp55xx: Initialize enable GPIO direction to output")
    Fixes: 92a81562e695 ("leds: lp55xx: Add multicolor framework support to lp55xx")
    Fixes: ac219bf3c9bd ("leds: lp55xx: Convert to use GPIO descriptors")
    Signed-off-by: Merlijn Wajer <merlijn@wizzup.org>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Pavel Machek <pavel@ucw.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 18e66ba7b3d617606c023e1ffe9777e556285797
Author: Sicelo A. Mhlongo <absicsz@gmail.com>
Date:   Sun Dec 12 23:40:06 2021 +0100

    ARM: dts: omap3-n900: Fix lp5523 for multi color

    [ Upstream commit e9af026a3b24f59d7af4609f73e0ef60a4d6d516 ]

    Since the LED multicolor framework support was added in commit
    92a81562e695 ("leds: lp55xx: Add multicolor framework support to lp55xx")
    LEDs on this platform stopped working.

    Fixes: 92a81562e695 ("leds: lp55xx: Add multicolor framework support to lp55xx")
    Fixes: ac219bf3c9bd ("leds: lp55xx: Convert to use GPIO descriptors")
    Signed-off-by: Merlijn Wajer <merlijn@wizzup.org>
    Signed-off-by: Sicelo A. Mhlongo <absicsz@gmail.com>
    Signed-off-by: Pavel Machek <pavel@ucw.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d2c375e9aeeedf90f127201f06db23048b5bd8f5
Author: Sudeep Holla <sudeep.holla@arm.com>
Date:   Thu Dec 9 09:21:46 2021 +0000

    mailbox: pcc: Handle all PCC subtypes correctly in pcc_mbox_irq

    [ Upstream commit 7215a7857e796c655ae1184b313556102fa8bc40 ]

    Commit c45ded7e1135 ("mailbox: pcc: Add support for PCCT extended PCC
    subspaces(type 3/4)") enabled the type3/4 of PCCT, but the change in
    pcc_mbox_irq breaks the other PCC subtypes.

    The kernel reports a warning on an Ampere eMag server

    -->8
     CPU: 0 PID: 0 Comm: swapper/0 Not tainted 5.16.0-rc4 #127
     Hardware name: MiTAC RAPTOR EV-883832-X3-0001/RAPTOR, BIOS 0.14 02/22/2019
     Call trace:
      dump_backtrace+0x0/0x200
      show_stack+0x20/0x30
      dump_stack_lvl+0x68/0x84
      dump_stack+0x18/0x34
      __report_bad_irq+0x54/0x17c
      note_interrupt+0x330/0x428
      handle_irq_event_percpu+0x90/0x98
      handle_irq_event+0x4c/0x148
      handle_fasteoi_irq+0xc4/0x188
      generic_handle_domain_irq+0x44/0x68
      gic_handle_irq+0x84/0x2ec
      call_on_irq_stack+0x28/0x34
      do_interrupt_handler+0x88/0x90
      el1_interrupt+0x48/0xb0
      el1h_64_irq_handler+0x18/0x28
      el1h_64_irq+0x7c/0x80

    Fixes: c45ded7e1135 ("mailbox: pcc: Add support for PCCT extended PCC subspaces(type
3/4)")
    Reported-by: Justin He <justin.he@arm.com>
    Tested-by: Justin He <justin.he@arm.com>
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b986b4ef835f6ef870ac152fb1ac9866330f94e4
Author: Sudeep Holla <sudeep.holla@arm.com>
Date:   Thu Dec 9 08:21:43 2021 +0000
```

```
    mailbox: pcc: Avoid using the uninitialized variable 'dev'

    [ Upstream commit 960c4056aadcf61983f8eaac159927a052f8cf01 ]

    Smatch static checker warns:

      |  drivers/mailbox/pcc.c:292 pcc_mbox_request_channel()
      |  error: uninitialized symbol 'dev'.

    Fix the same by using pr_err instead of dev_err as the variable 'dev'
    is uninitialized at that stage.

    Fixes: ce028702ddbc ("mailbox: pcc: Move bulk of PCCT parsing into pcc_mbox_probe")
    Cc: Jassi Brar <jassisinghbrar@gmail.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3309cb57ae7110ee2bd04642ce3c08b3eb035817
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Wed Nov 24 17:51:26 2021 +0300

    mailbox: imx: Fix an IS_ERR() vs NULL bug

    [ Upstream commit 05d06f37196b2e3abeff2b98b785c8803865e646 ]

    The devm_kzalloc() function does not return error pointers, it returns
    NULL on failure.

    Fixes: 97961f78e8bc ("mailbox: imx: support i.MX8ULP S4 MU")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f27ec926fd66d2c937f351338d535d818388df5
Author: jason-jh.lin <jason-jh.lin@mediatek.com>
Date:   Thu Dec 23 22:51:55 2021 +0800

    mailbox: fix gce_num of mt8192 driver data

    [ Upstream commit 35ca43710f792ce183312fdc7e4b2bb0b721a173 ]

    Because mt8192 only have 1 gce, the gce_num should be 1.

    Fixes: 85dfdbfc13ea ("mailbox: cmdq: add multi-gce clocks support for mt8195")
    Signed-off-by: jason-jh.lin <jason-jh.lin@mediatek.com>
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Jassi Brar <jaswinder.singh@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06649e77f05dd23253d10687ba3292cf78a11458
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Fri Jan 7 18:57:22 2022 +0000

    MIPS: compressed: Fix build with ZSTD compression

    [ Upstream commit c5c7440fe7f74645940d5c9e2c49cd7efb706a4f ]

    Fix the following build issues:

    mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
`FSE_buildDTable_internal':
```

```
      decompress.c:(.text.FSE_buildDTable_internal+0x2cc): undefined reference to `__clzdi2'
        mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `BIT_initDStream':
        decompress.c:(.text.BIT_initDStream+0x7c): undefined reference to `__clzdi2'
        mips64el-linux-ld: decompress.c:(.text.BIT_initDStream+0x158): undefined reference to
  `__clzdi2'
        mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `ZSTD_buildFSETable_body_default.constprop.0':
      decompress.c:(.text.ZSTD_buildFSETable_body_default.constprop.0+0x2a8): undefined
  reference to `__clzdi2'
        mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `FSE_readNCount_body_default':
      decompress.c:(.text.FSE_readNCount_body_default+0x130): undefined reference to
  `__ctzdi2'
      mips64el-linux-ld: decompress.c:(.text.FSE_readNCount_body_default+0x1a4): undefined
  reference to `__ctzdi2'
      mips64el-linux-ld: decompress.c:(.text.FSE_readNCount_body_default+0x2e4): undefined
  reference to `__clzdi2'
        mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `HUF_readStats_body_default':
      decompress.c:(.text.HUF_readStats_body_default+0x184): undefined reference to
  `__clzdi2'
      mips64el-linux-ld: decompress.c:(.text.HUF_readStats_body_default+0x1b4): undefined
  reference to `__clzdi2'
        mips64el-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `ZSTD_DCtx_getParameter':
      decompress.c:(.text.ZSTD_DCtx_getParameter+0x60): undefined reference to `__clzdi2'

    Fixes: a510b616131f ("MIPS: Add support for ZSTD-compressed kernels")
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Nick Terrell <terrelln@fb.com>
    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 9ea3913af88d992c03dc4b0e5819ce8937e0af36
Author: Stephen Boyd <swboyd@chromium.org>
Date:   Fri Jan 7 11:42:32 2022 -0800

    of/fdt: Don't worry about non-memory region overlap for no-map

    [ Upstream commit da17d6905d29ddcdc04b2fdc37ed8cf1e8437cc8 ]

    In commit 8a5a75e5e9e5 ("of/fdt: Make sure no-map does not remove
    already reserved regions") we returned -EBUSY when trying to mark
    regions as no-map when they intersect with reserved memory. The goal was
    to find bad no-map reserved memory DT nodes that would unmap the kernel
    text/data sections.

    The problem is the reserved memory check will still trigger if the DT
    has a /memreserve/ that completely subsumes the no-map memory carveouts
    in the reserved memory node _and_ that region is also not part of the
    memory reg property. For example in sc7180.dtsi we have the following
    reserved-memory and memory node:

          memory@80000000 {
              /* We expect the bootloader to fill in the size */
              reg = <0 0x80000000 0 0>;
          };

          smem_mem: memory@80900000 {
                  reg = <0x0 0x80900000 0x0 0x200000>;
                  no-map;
          };
```

```
    and the memreserve filled in by the bootloader is

          /memreserve/ 0x80800000 0x400000;

    while the /memory node is transformed into

          memory@80000000 {
              /* The bootloader fills in the size, and adds another region */
              reg = <0 0x80000000 0 0x00800000>,
                    <0 0x80c00000 0 0x7f200000>;
          };

    The smem region is doubly reserved via /memreserve/ and by not being
    part of the /memory reg property. This leads to the following warning
    printed at boot.

     OF: fdt: Reserved memory: failed to reserve memory for node 'memory@80900000': base
0x0000000080900000, size 2 MiB

    Otherwise nothing really goes wrong because the smem region is not going
    to be mapped by the kernel's direct linear mapping given that it isn't
    part of the memory node. Therefore, let's only consider this to be a
    problem if we're trying to mark a region as no-map and it is actually
    memory that we're intending to keep out of the kernel's direct mapping
    but it's already been reserved.

    Acked-by: Mike Rapoport <rppt@kernel.org>
    Cc: Douglas Anderson <dianders@chromium.org>
    Cc: Nicolas Boichat <drinkcat@chromium.org>
    Cc: Quentin Perret <qperret@google.com>
    Cc: Jan Kiszka <jan.kiszka@siemens.com>
    Fixes: 8a5a75e5e9e5 ("of/fdt: Make sure no-map does not remove already reserved
regions")
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20220107194233.2793146-1-swboyd@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1180c4d60c6506adf251ec5c05b8b50c2bb276cd
Author: Baruch Siach <baruch@tkos.co.il>
Date:   Thu Dec 30 18:31:52 2021 +0200

    of: base: Fix phandle argument length mismatch error message

    [ Upstream commit 94a4950a4acff39b5847cc1fee4f65e160813493 ]

    The cell_count field of of_phandle_iterator is the number of cells we
    expect in the phandle arguments list when cells_name is missing. The
    error message should show the number of cells we actually see.

    Fixes: af3be70a3211 ("of: Improve of_phandle_iterator_next() error message")
    Cc: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Baruch Siach <baruch@tkos.co.il>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Link:
https://lore.kernel.org/r/96519ac55be90a63fa44afe01480c30d08535465.1640881913.git.baruch@tko
s.co.il
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06cceda7eff54334b2c3909b60b4c4c91cc4f434
Author: Conor Dooley <conor.dooley@microchip.com>
Date:   Thu Dec 23 15:42:44 2021 +0000

    clk: bm1880: remove kfrees on static allocations
```

```
    [ Upstream commit c861c1be3897845313a0df47804b1db37c7052e1 ]

    bm1880_clk_unregister_pll & bm1880_clk_unregister_div both try to
    free statically allocated variables, so remove those kfrees.

    For example, if we take L703 kfree(div_hw):
    - div_hw is a bm1880_div_hw_clock pointer
    - in bm1880_clk_register_plls this is pointed to an element of arg1:
      struct bm1880_div_hw_clock *clks
    - in the probe, where bm1880_clk_register_plls is called arg1 is
      bm1880_div_clks, defined on L371:
      static struct bm1880_div_hw_clock bm1880_div_clks[]

    Signed-off-by: Conor Dooley <conor.dooley@microchip.com>
    Fixes: 1ab4601da55b ("clk: Add common clock driver for BM1880 SoC")
    Link: https://lore.kernel.org/r/20211223154244.1024062-1-conor.dooley@microchip.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 44decfc75a461b6dbf367f18f3899a0749c68d34
Author: Shengjiu Wang <shengjiu.wang@nxp.com>
Date:   Wed Jan 5 19:08:03 2022 +0800

    ASoC: fsl_asrc: refine the check of available clock divider

    [ Upstream commit 320386343451ab6a3577e0ee200dac56a6182944 ]

    According to RM, the clock divider range is from 1 to 8, clock
    prescaling ratio may be any power of 2 from 1 to 128.
    So the supported divider is not all the value between
    1 and 1024, just limited value in that range.

    Create table for the supported divder and add function to
    check the clock divider is available by comparing with
    the table.

    Fixes: d0250cf4f2ab ("ASoC: fsl_asrc: Add an option to select internal ratio mode")
    Signed-off-by: Shengjiu Wang <shengjiu.wang@nxp.com>
    Link: https://lore.kernel.org/r/1641380883-20709-1-git-send-email-shengjiu.wang@nxp.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e496eec9c330f9a80f4a76871582b32547fda080
Author: Shengjiu Wang <shengjiu.wang@nxp.com>
Date:   Tue Jan 4 18:40:35 2022 +0800

    ASoC: imx-card: improve the sound quality for low rate

    [ Upstream commit 3969341813eb56d2dfc39bb64229359a6ae3c195 ]

    According to RM, on auto mode:
    For codec AK4458 and AK4497, the lowest ratio of MLCK/FS is 256
    if sample rate is 8kHz-48kHz,
    For codec AK5558, the lowest ratio of MLCK/FS is 512 if sample
    rate is 8kHz-48kHz.

    With these setting the sound quality for 8kHz-48kHz can be improved.

    Fixes: aa736700f42f ("ASoC: imx-card: Add imx-card machine driver")
    Signed-off-by: Shengjiu Wang <shengjiu.wang@nxp.com>
    Link: https://lore.kernel.org/r/1641292835-19085-4-git-send-email-shengjiu.wang@nxp.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 639ae7e27bacd26da10ee80b0fc698de8df1ac09
```

```
Author: Shengjiu Wang <shengjiu.wang@nxp.com>
Date:   Tue Jan 4 18:40:34 2022 +0800

    ASoC: imx-card: Fix mclk calculation issue for akcodec

    [ Upstream commit f331ae5fa59fbfb748317b290648fc3f1a50d932 ]

    Transfer the refined slots and slot_width to akcodec_get_mclk_rate()
    for mclk calculation, otherwise the mclk frequency does not match
    with the slots and slot_width for S16_LE format, because the default
    slot_width is 32.

    Fixes: aa736700f42f ("ASoC: imx-card: Add imx-card machine driver")
    Signed-off-by: Shengjiu Wang <shengjiu.wang@nxp.com>
    Link: https://lore.kernel.org/r/1641292835-19085-3-git-send-email-shengjiu.wang@nxp.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 360511ef065b93bdbc615b084528c390b9b2d9d6
Author: Shengjiu Wang <shengjiu.wang@nxp.com>
Date:   Tue Jan 4 18:40:33 2022 +0800

    ASoC: imx-card: Need special setting for ak4497 on i.MX8MQ

    [ Upstream commit 3349b3d0c63b8b6fcca58156d72407f0b2e101ac ]

    The SAI on i.MX8MQ don't support one2one ratio for mclk:bclk, so
    the mclk frequency exceeds the supported range of codec for
    the case that sample rate is larger than 705kHZ and format is
    S32_LE. Update the supported width for such case.

    Fixes: aa736700f42f ("ASoC: imx-card: Add imx-card machine driver")
    Signed-off-by: Shengjiu Wang <shengjiu.wang@nxp.com>
    Link: https://lore.kernel.org/r/1641292835-19085-2-git-send-email-shengjiu.wang@nxp.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c8ca01dc3e65fed7933d2ca91b457ab027459960
Author: Taniya Das <tdas@codeaurora.org>
Date:   Mon Dec 20 22:13:56 2021 +0530

    clk: qcom: gcc-sc7280: Mark gcc_cfg_noc_lpass_clk always enabled

    [ Upstream commit 9c337073d9d81a145434b22f42dc3128ecd17730 ]

    The gcc cfg noc lpass clock is required to be always enabled for the
    LPASS core and audio drivers to be functional.

    Fixes: a3cc092196ef ("clk: qcom: Add Global Clock controller (GCC) driver for SC7280")
    Signed-off-by: Taniya Das <tdas@codeaurora.org>
    Link: https://lore.kernel.org/r/1640018638-19436-4-git-send-email-tdas@codeaurora.org
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05413dcadb28cebb6aed12b44139439ce8086bae
Author: Kamal Heib <kamalheib1@gmail.com>
Date:   Mon Dec 20 17:25:30 2021 +0200

    RDMA/cxgb4: Set queue pair state when being queried

    [ Upstream commit e375b9c92985e409c4bb95dd43d34915ea7f5e28 ]

    The API for ib_query_qp requires the driver to set cur_qp_state on return,
    add the missing set.
```

```
    Fixes: 67bbc05512d8 ("RDMA/cxgb4: Add query_qp support")
    Link: https://lore.kernel.org/r/20211220152530.60399-1-kamalheib1@gmail.com
    Signed-off-by: Kamal Heib <kamalheib1@gmail.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a0ecfd10d669c99fb04b1a973a5da8cd4e4fcc41
Author: Huang Pei <huangpei@loongson.cn>
Date:   Wed Dec 15 16:44:57 2021 +0800

    MIPS: fix local_{add,sub}_return on MIPS64

    [ Upstream commit 277c8cb3e8ac199f075bf9576ad286687ed17173 ]

    Use "daddu/dsubu" for long int on MIPS64 instead of "addu/subu"

    Fixes: 7232311ef14c ("local_t: mips extension")
    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 80cf54e282b25b1b4496c6cdd4237d095e323115
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Fri Dec 10 14:36:27 2021 -0700

    cxl/core: Remove cxld_const_init in cxl_decoder_alloc()

    [ Upstream commit be185c2988b48db65348d94168c793bdbc8d23c3 ]

    Commit 48667f676189 ("cxl/core: Split decoder setup into alloc + add")
    aimed to fix a large stack frame warning but from v5 to v6, it
    introduced a new instance of the warning due to allocating
    cxld_const_init on the stack, which was done due to the use of const on
    the nr_target member of the cxl_decoder struct. With ARCH=arm
    allmodconfig minus CONFIG_KASAN:

    GCC 11.2.0:

    drivers/cxl/core/bus.c: In function ‘cxl_decoder_alloc’:
    drivers/cxl/core/bus.c:523:1: error: the frame size of 1032 bytes is larger than 1024
bytes [-Werror=frame-larger-than=]
      523 | }
          | ^
    cc1: all warnings being treated as errors

    Clang 12.0.1:

    drivers/cxl/core/bus.c:486:21: error: stack frame size of 1056 bytes in function
'cxl_decoder_alloc' [-Werror,-Wframe-larger-than=]
    struct cxl_decoder *cxl_decoder_alloc(struct cxl_port *port, int nr_targets)
                        ^
    1 error generated.

    Revert that part of the change, which makes the stack frame of
    cxl_decoder_alloc() much more reasonable.

    Fixes: 48667f676189 ("cxl/core: Split decoder setup into alloc + add")
    Link: https://github.com/ClangBuiltLinux/linux/issues/1539
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20211210213627.2477370-1-nathan@kernel.org
    Signed-off-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit abc38a36dfaef272a0822b742595450f4aee3bbd
```

```
Author: Christian A. Ehrhardt <lk@c--e.de>
Date:   Fri Dec 31 14:44:32 2021 +0100

    ALSA: hda/cs8409: Fix Jack detection after resume

    [ Upstream commit 57f234248ff925d88caedf4019ec84e6ecb83909 ]

    The suspend code unconditionally sets ->hp_jack_in and ->mic_jack_in
    to zero but without reporting this status change to the HDA core.
    To compensate for this, always assume a status change on the
    first unsol event after boot or resume.

    Fixes: 424e531b47f8 ("ALSA: hda/cs8409: Ensure Type Detection is only run on startup
when necessary")
    Signed-off-by: Christian A. Ehrhardt <lk@c--e.de>
    Link: https://lore.kernel.org/r/20211231134432.atwmuzeceqiklcoa@cae.in-ulm.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b9f0e38e782a6c5ecb34314201a6128892547164
Author: Christian A. Ehrhardt <lk@c--e.de>
Date:   Fri Dec 31 14:12:21 2021 +0100

    ALSA: hda/cs8409: Increase delay during jack detection

    [ Upstream commit 8cd07657177006b67cc1610e4466cc75ad781c05 ]

    Commit c8b4f0865e82 reduced delays related to cs42l42 jack
    detection. However, the change was too aggressive. As a result
    internal speakers on DELL Inspirion 3501 are not detected.

    Increase the delay in cs42l42_run_jack_detect() a bit.

    Fixes: c8b4f0865e82 ("ALSA: hda/cs8409: Remove unnecessary delays")
    Signed-off-by: Christian A. Ehrhardt <lk@c--e.de>
    Link: https://lore.kernel.org/r/20211231131221.itwotyfk5qomn7n6@cae.in-ulm.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bd68d7063cf6f8a7e42621652e27045de5b8134
Author: Alyssa Ross <hi@alyssa.is>
Date:   Tue Jan 4 13:22:16 2022 +0000

    ASoC: fsl_mqs: fix MODULE_ALIAS

    [ Upstream commit 9f3d45318dd9e739ed62e4218839a7a824d3cced ]

    modprobe can't handle spaces in aliases.

    Fixes: 9e28f6532c61 ("ASoC: fsl_mqs: Add MQS component driver")
    Signed-off-by: Alyssa Ross <hi@alyssa.is>
    Link: https://lore.kernel.org/r/20220104132218.1690103-1-hi@alyssa.is
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1400d9bbd1b2e685717dcbd03ea63428d3d48f3c
Author: Ammar Faizi <ammarfaizi2@gmail.com>
Date:   Sun Dec 26 20:54:02 2021 +0700

    powerpc/xive: Add missing null check after calling kmalloc

    [ Upstream commit 18dbfcdedc802f9500b2c29794f22a31d27639c0 ]

    Commit 930914b7d528fc ("powerpc/xive: Add a debugfs file to dump
    internal XIVE state") forgot to add a null check.
```

```
    Add it.

    Fixes: 930914b7d528fc6b0249bffc00564100bcf6ef75 ("powerpc/xive: Add a debugfs file to
dump internal XIVE state")
    Signed-off-by: Ammar Faizi <ammarfaizi2@gmail.com>
    Reviewed-by: Cédric Le Goater <clg@kaod.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20211226135314.251221-1-ammar.faizi@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f68114e254128aebe053851f37f9c940b791e971
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Tue Dec 28 16:05:53 2021 -0800

    mips: bcm63xx: add support for clk_set_parent()

    [ Upstream commit 6f03055d508ff4feb8db02ba3df9303a1db8d381 ]

    The MIPS BMC63XX subarch does not provide/support clk_set_parent().
    This causes build errors in a few drivers, so add a simple implementation
    of that function so that callers of it will build without errors.

    Fixes these build errors:

    ERROR: modpost: "clk_set_parent" [sound/soc/jz4740/snd-soc-jz4740-i2s.ko] undefined!
    ERROR: modpost: "clk_set_parent" [sound/soc/atmel/snd-soc-atmel-i2s.ko] undefined!

    Fixes: e7300d04bd08 ("MIPS: BCM63xx: Add support for the Broadcom BCM63xx family of
SOCs." )
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Reviewed-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 54e2a610dfeeb6bbf63c4f47bf9a682abfcb90b2
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Tue Dec 28 16:03:45 2021 -0800

    mips: lantiq: add support for clk_set_parent()

    [ Upstream commit 76f66dfd60dc5d2f9dec22d99091fea1035c5d03 ]

    Provide a simple implementation of clk_set_parent() in the lantiq
    subarch so that callers of it will build without errors.

    Fixes these build errors:

    ERROR: modpost: "clk_set_parent" [sound/soc/jz4740/snd-soc-jz4740-i2s.ko] undefined!
    ERROR: modpost: "clk_set_parent" [sound/soc/atmel/snd-soc-atmel-i2s.ko] undefined!

    Fixes: 171bb2f19ed6 ("MIPS: Lantiq: Add initial support for Lantiq SoCs")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Reported-by: kernel test robot <lkp@intel.com>
    --to=linux-mips@vger.kernel.org --cc="John Crispin <john@phrozen.org>" --cc="Jonathan
Cameron <jic23@kernel.org>" --cc="Russell King <linux@armlinux.org.uk>" --cc="Andy
Shevchenko <andy.shevchenko@gmail.com>" --cc=alsa-devel@alsa-project.org --to="Thomas
Bogendoerfer <tsbogend@alpha.franken.de>"
    Reviewed-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 955a5794c47de18a57d93a9726b4666c47b9d850
Author: Sameer Pujar <spujar@nvidia.com>
```

```
Date:   Thu Dec 23 17:23:51 2021 +0530

    arm64: tegra: Remove non existent Tegra194 reset

    [ Upstream commit 146b3a77af8091cabbd1decc51d67799e69682d2 ]

    Tegra194 does not really have "hda2codec_2x" related reset. Hence drop
    this entry to reflect actual HW.

    Fixes: 4878cc0c9fab ("arm64: tegra: Add HDA controller on Tegra194")
    Signed-off-by: Sameer Pujar <spujar@nvidia.com>
    Link: https://lore.kernel.org/r/1640260431-11613-4-git-send-email-spujar@nvidia.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fb5ec0543db5d09ce70fec39c343d5d22519abe2
Author: Trevor Wu <trevor.wu@mediatek.com>
Date:   Thu Dec 30 16:47:30 2021 +0800

    ASoC: mediatek: mt8195: correct pcmif BE dai control flow

    [ Upstream commit 2355028c0c54c03afb66c589347f1dc9f6fe2e38 ]

    Originally, the conditions for preventing reentry are not correct.
    dai->component->active is not the state specifically for pcmif dai, so it
    is not a correct condition to indicate the status of pcmif dai.
    On the other hand, snd_soc_dai_stream_actvie() in prepare ops for both
    playback and capture possibly return true at the first entry when these
    two streams are opened at the same time.

    In the patch, I refer to the implementation in mt8192-dai-pcm.c.
    Clock and enabling bit for PCMIF are managed by DAPM, and the condition
    for prepare ops is replaced by the status of dai widget.

    Fixes: 1f95c019115c ("ASoC: mediatek: mt8195: support pcm in platform driver")
    Signed-off-by: Trevor Wu <trevor.wu@mediatek.com>
    Link: https://lore.kernel.org/r/20211230084731.31372-2-trevor.wu@mediatek.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d8e8b1a2cb8e0b4fc6cd6ddf7cf290fc19eeaefd
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Tue Dec 28 12:55:22 2021 +0000

    misc: lattice-ecp3-config: Fix task hung when firmware load failed

    [ Upstream commit fcee5ce50bdb21116711e38635e3865594af907e ]

    When firmware load failed, kernel report task hung as follows:

    INFO: task xrun:5191 blocked for more than 147 seconds.
          Tainted: G        W         5.16.0-rc5-next-20211220+ #11
    "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
    task:xrun            state:D stack:    0 pid: 5191 ppid:   270 flags:0x00000004
    Call Trace:
     __schedule+0xc12/0x4b50 kernel/sched/core.c:4986
     schedule+0xd7/0x260 kernel/sched/core.c:6369 (discriminator 1)
     schedule_timeout+0x7aa/0xa80 kernel/time/timer.c:1857
     wait_for_completion+0x181/0x290 kernel/sched/completion.c:85
     lattice_ecp3_remove+0x32/0x40 drivers/misc/lattice-ecp3-config.c:221
     spi_remove+0x72/0xb0 drivers/spi/spi.c:409

    lattice_ecp3_remove() wait for signals from firmware loading, but when
    load failed, firmware_load() does not send this signal. This cause
    device remove hung. Fix it by sending signal even if load failed.
```

```
    Fixes: 781551df57c7 ("misc: Add Lattice ECP3 FPGA configuration via SPI")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Link: https://lore.kernel.org/r/20211228125522.3122284-1-weiyongjun1@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5fcdc740077c9cbfcad8978d1d276b0d28a77213
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Dec 28 11:40:26 2021 +0800

    ASoC: samsung: idma: Check of ioremap return value

    [ Upstream commit 3ecb46755eb85456b459a1a9f952c52986bce8ec ]

    Because of the potential failure of the ioremap(), the buf->area could
    be NULL.
    Therefore, we need to check it and return -ENOMEM in order to transfer
    the error.

    Fixes: f09aecd50f39 ("ASoC: SAMSUNG: Add I2S0 internal dma driver")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20211228034026.1659385-1-jiasheng@iscas.ac.cn
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab87117cddddb3550d2ed417df4a003dbfdea97e
Author: Tom Rix <trix@redhat.com>
Date:   Fri Dec 24 07:08:33 2021 -0800

    iio: chemical: sunrise_co2: set val parameter only on success

    [ Upstream commit 38ac2f038666521f94d4fa37b5a9441cef832ccf ]

    Clang static analysis reports this representative warning

    sunrise_co2.c:410:9: warning: Assigned value is garbage or undefined
      *val = value;
           ^ ~~~~~

    The ealier call to sunrise_read_word can fail without setting
    value.  So defer setting val until we know the read was successful.

    Fixes: c397894e24f1 ("iio: chemical: Add Senseair Sunrise 006-0-007 driver")
    Signed-off-by: Tom Rix <trix@redhat.com>
    Link: https://lore.kernel.org/r/20211224150833.3278236-1-trix@redhat.com
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41bd13696a9fec1d5b438755012abf50d3d19cda
Author: Swapnil Jakhade <sjakhade@cadence.com>
Date:   Thu Dec 23 07:01:33 2021 +0100

    phy: cadence: Sierra: Fix to get correct parent for mux clocks

    [ Upstream commit da08aab940092a050a4fb2857ed9479d2b0e03c4 ]

    Fix get_parent() callback to return the correct index of the parent for
    PLL_CMNLC1 clock. Add a separate table of register values corresponding
    to the parent index for PLL_CMNLC1. Update set_parent() callback
    accordingly.

    Fixes: 28081b72859f ("phy: cadence: Sierra: Model PLL_CMNLC and PLL_CMNLC1 as clocks
```

```
  (mux clocks)")
      Signed-off-by: Swapnil Jakhade <sjakhade@cadence.com>
      Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
      Link: https://lore.kernel.org/r/20211223060137.9252-12-sjakhade@cadence.com
      Signed-off-by: Vinod Koul <vkoul@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 949255b0a0dcc6178378b31b79c6022b767cb13b
  Author: Michael Ellerman <mpe@ellerman.id.au>
  Date:   Wed Dec 22 00:51:00 2021 +1100

      powerpc/64s: Use EMIT_WARN_ENTRY for SRR debug warnings

      [ Upstream commit fd1eaaaaa6864b5fb8f99880fcefb49760b8fe4e ]

      When CONFIG_PPC_RFI_SRR_DEBUG=y we check the SRR values before returning
      from interrupts. This is done in asm using EMIT_BUG_ENTRY, and passing
      BUGFLAG_WARNING.

      However that fails to create an exception table entry for the warning,
      and so do_program_check() fails the exception table search and proceeds
      to call _exception(), resulting in an oops like:

        Oops: Exception in kernel mode, sig: 5 [#1]
        LE PAGE_SIZE=64K MMU=Radix SMP NR_CPUS=2048 NUMA pSeries
        Modules linked in:
        CPU: 2 PID: 1204 Comm: sigreturn_unali Tainted: P                   5.16.0-rc2-00194-
  g91ca3d4f77c5 #12
        NIP:  c00000000000c5b0 LR: 0000000000000000 CTR: 0000000000000000
        ...
        NIP [c00000000000c5b0] system_call_common+0x150/0x268
        LR [0000000000000000] 0x0
        Call Trace:
        [c00000000db73e10] [c00000000000c558] system_call_common+0xf8/0x268 (unreliable)
        ...
        Instruction dump:
        7cc803a6 888d0931 2c240000 4082001c 38800000 988d0931 e8810170 e8a10178
        7c9a03a6 7cbb03a6 7d7a02a6 e9810170 <7f0b6088> 7d7b02a6 e9810178 7f0b6088

      We should instead use EMIT_WARN_ENTRY, which creates an exception table
      entry for the warning, allowing the warning to be correctly recognised,
      and the code to resume after printing the warning.

      Note however that because this warning is buried deep in the interrupt
      return path, we are not able to recover from it (due to MSR_RI being
      clear), so we still end up in die() with an unrecoverable exception.

      Fixes: 59dc5bfca0cb ("powerpc/64s: avoid reloading (H)SRR registers if they are still
  valid")
      Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
      Link: https://lore.kernel.org/r/20211221135101.2085547-2-mpe@ellerman.id.au
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit e96874c8d1a39f032d1886d4e71d5d4959459cba
  Author: Michael Ellerman <mpe@ellerman.id.au>
  Date:   Wed Dec 22 00:50:59 2021 +1100

      powerpc/64s: Mask NIP before checking against SRR0

      [ Upstream commit 314f6c23dd8d417281eb9e8a516dd98036f2e7b3 ]

      When CONFIG_PPC_RFI_SRR_DEBUG=y we check that NIP and SRR0 match when
      returning from interrupts. This can trigger falsely if NIP has either of
      its two low bits set via sigreturn or ptrace, while SRR0 has its low two
      bits masked in hardware.
```

```
    As a quick fix make sure to mask the low bits before doing the check.

    Fixes: 59dc5bfca0cb ("powerpc/64s: avoid reloading (H)SRR registers if they are still
valid")
    Reported-by: Sachin Sant <sachinp@linux.vnet.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Tested-by: Sachin Sant <sachinp@linux.vnet.ibm.com>
    Link: https://lore.kernel.org/r/20211221135101.2085547-1-mpe@ellerman.id.au
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0bf3c9e05ca6837ac756ec7d9de70b44603da12
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Dec 22 09:51:57 2021 +0800

    ASoC: mediatek: Check for error clk pointer

    [ Upstream commit 9de2b9286a6dd16966959b3cb34fc2ddfd39213e ]

    Yes, you are right and now the return code depending on the
    init_clks().

    Fixes: 6078c651947a ("soc: mediatek: Refine scpsys to support multiple platform")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Link: https://lore.kernel.org/r/20211222015157.1025853-1-jiasheng@iscas.ac.cn
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ba1c03537a2f7aa5268d8c79c521ecbb4b56fe3
Author: Ryuta NAKANISHI <nakanishi.ryuta@socionext.com>
Date:   Wed Dec 22 14:19:29 2021 +0900

    phy: uniphier-usb3ss: fix unintended writing zeros to PHY register

    [ Upstream commit 898c7a9ec81620125f2463714a0f4dea18ad6e54 ]

    Similar to commit 4a90bbb478db ("phy: uniphier-pcie: Fix updating phy
    parameters"), in function uniphier_u3ssphy_set_param(), unintentionally
    write zeros to other fields when writing PHY registers.

    Fixes: 5ab43d0f8697 ("phy: socionext: add USB3 PHY driver for UniPhier SoC")
    Signed-off-by: Ryuta NAKANISHI <nakanishi.ryuta@socionext.com>
    Signed-off-by: Kunihiko Hayashi <hayashi.kunihiko@socionext.com>
    Link: https://lore.kernel.org/r/1640150369-4134-1-git-send-email-
hayashi.kunihiko@socionext.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7369b31ee7cd095dba8993a9a624c955bd390c43
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Mon Dec 20 19:21:26 2021 +0800

    scsi: block: pm: Always set request queue runtime active in blk_post_runtime_resume()

    [ Upstream commit 6e1fcab00a23f7fe9f4fe9704905a790efa1eeab ]

    John Garry reported a deadlock that occurs when trying to access a
    runtime-suspended SATA device.  For obscure reasons, the rescan procedure
    causes the link to be hard-reset, which disconnects the device.

    The rescan tries to carry out a runtime resume when accessing the device.
    scsi_rescan_device() holds the SCSI device lock and won't release it until
    it can put commands onto the device's block queue.  This can't happen until
    the queue is successfully runtime-resumed or the device is unregistered.
    But the runtime resume fails because the device is disconnected, and
```

```
    __scsi_remove_device() can't do the unregistration because it can't get the
    device lock.

    The best way to resolve this deadlock appears to be to allow the block
    queue to start running again even after an unsuccessful runtime resume.
    The idea is that the driver or the SCSI error handler will need to be able
    to use the queue to resolve the runtime resume failure.

    This patch removes the err argument to blk_post_runtime_resume() and makes
    the routine act as though the resume was successful always.  This fixes the
    deadlock.

    Link: https://lore.kernel.org/r/1639999298-244569-4-git-send-email-
chenxiang66@hisilicon.com
    Fixes: e27829dc92e5 ("scsi: serialize ->rescan against ->remove")
    Reported-and-tested-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Signed-off-by: Xiang Chen <chenxiang66@hisilicon.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec00c3715d29cbc82fa25ba68aef6a7ee46c518a
Author: Pingfan Liu <kernelfans@gmail.com>
Date:   Wed Dec 15 10:13:48 2021 +0800

    efi: apply memblock cap after memblock_add()

    [ Upstream commit b398123bff3bcbc1facb0f29bf6e7b9f1bc55931 ]

    On arm64, during kdump kernel saves vmcore, it runs into the following bug:
    ...
    [   15.148919] usercopy: Kernel memory exposure attempt detected from SLUB object
'kmem_cache_node' (offset 0, size 4096)!
    [   15.159707] ------------[ cut here ]------------
    [   15.164311] kernel BUG at mm/usercopy.c:99!
    [   15.168482] Internal error: Oops - BUG: 0 [#1] SMP
    [   15.173261] Modules linked in: xfs libcrc32c crct10dif_ce ghash_ce sha2_ce
sha256_arm64 sha1_ce sbsa_gwdt ast i2c_algo_bit drm_vram_helper drm_kms_helper syscopyarea
sysfillrect sysimgblt fb_sys_fops cec drm_ttm_helper ttm drm nvme nvme_core xgene_hwmon
i2c_designware_platform i2c_designware_core dm_mirror dm_region_hash dm_log dm_mod overlay
squashfs zstd_decompress loop
    [   15.206186] CPU: 0 PID: 542 Comm: cp Not tainted 5.16.0-rc4 #1
    [   15.212006] Hardware name: GIGABYTE R272-P30-JG/MP32-AR0-JG, BIOS F12 (SCP:
1.5.20210426) 05/13/2021
    [   15.221125] pstate: 60400009 (nZCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--)
    [   15.228073] pc : usercopy_abort+0x9c/0xa0
    [   15.232074] lr : usercopy_abort+0x9c/0xa0
    [   15.236070] sp : ffff8000121abba0
    [   15.239371] x29: ffff8000121abbb0 x28: 0000000000003000 x27: 0000000000000000
    [   15.246494] x26: 0000000080000400 x25: 0000ffff885c7000 x24: 0000000000000000
    [   15.253617] x23: 000007ff80400000 x22: ffff07ff80401000 x21: 0000000000000001
    [   15.260739] x20: 0000000000001000 x19: ffff07ff80400000 x18: ffffffffffffffff
    [   15.267861] x17: 656a626f2042554c x16: 53206d6f72662064 x15: 6574636574656420
    [   15.274983] x14: 74706d6574746120 x13: 2129363930342065 x12: 7a6973202c302074
    [   15.282105] x11: ffffc8b041d1b148 x10: 00000000ffff8000 x9 : ffffc8b04012812c
    [   15.289228] x8 : 00000000ffff7fff x7 : ffffc8b041d1b148 x6 : 0000000000000000
    [   15.296349] x5 : 0000000000000000 x4 : 0000000000007fff x3 : 0000000000000000
    [   15.303471] x2 : 0000000000000000 x1 : ffff07ff8c064800 x0 : 000000000000006b
    [   15.310593] Call trace:
    [   15.313027]  usercopy_abort+0x9c/0xa0
    [   15.316677]  __check_heap_object+0xd4/0xf0
    [   15.320762]  __check_object_size.part.0+0x160/0x1e0
    [   15.325628]  __check_object_size+0x2c/0x40
    [   15.329711]  copy_oldmem_page+0x7c/0x140
```

```
[   15.333623]  read_from_oldmem.part.0+0xfc/0x1c0
[   15.338142]  __read_vmcore.constprop.0+0x23c/0x350
[   15.342920]  read_vmcore+0x28/0x34
[   15.346309]  proc_reg_read+0xb4/0xf0
[   15.349871]  vfs_read+0xb8/0x1f0
[   15.353088]  ksys_read+0x74/0x100
[   15.356390]  __arm64_sys_read+0x28/0x34
...
```

This bug introduced by commit b261dba2fdb2 ("arm64: kdump: Remove custom linux,usable-memory-range handling"), which moves memblock_cap_memory_range() to fdt, but it breaches the rules that memblock_cap_memory_range() should come after memblock_add() etc as said in commit e888fa7bb882 ("memblock: Check memory add/cap ordering").

As a consequence, the virtual address set up by copy_oldmem_page() does not bail out from the test of virt_addr_valid() in check_heap_object(), and finally hits the BUG_ON().

Since memblock allocator has no idea about when the memblock is fully populated, while efi_init() is aware, so tackling this issue by calling the interface early_init_dt_check_for_usable_mem_range() exposed by of/fdt.

Fixes: b261dba2fdb2 ("arm64: kdump: Remove custom linux,usable-memory-range handling")
Signed-off-by: Pingfan Liu <kernelfans@gmail.com>
Cc: Rob Herring <robh+dt@kernel.org>
Cc: Zhen Lei <thunder.leizhen@huawei.com>
Cc: Catalin Marinas <catalin.marinas@arm.com>
Cc: Will Deacon <will@kernel.org>
Cc: Andrew Morton <akpm@linux-foundation.org>
Cc: Mike Rapoport <rppt@kernel.org>
Cc: Geert Uytterhoeven <geert+renesas@glider.be>
Cc: Frank Rowand <frowand.list@gmail.com>
Cc: Ard Biesheuvel <ardb@kernel.org>
Cc: Nick Terrell <terrelln@fb.com>
Cc: linux-arm-kernel@lists.infradead.org
To: devicetree@vger.kernel.org
To: linux-efi@vger.kernel.org
Acked-by: Ard Biesheuvel <ardb@kernel.org>
Signed-off-by: Rob Herring <robh@kernel.org>
Link: https://lore.kernel.org/r/20211215021348.8766-1-kernelfans@gmail.com
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17c8eead5563fb6e8a789f2d31b7d08d9f8a3408
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date: Tue Dec 14 12:01:56 2021 +0800

of: fdt: Aggregate the processing of "linux,usable-memory-range"

[Upstream commit 8347b41748c3019157312fbe7f8a6792ae396eb7]

Currently, we parse the "linux,usable-memory-range" property in early_init_dt_scan_chosen(), to obtain the specified memory range of the crash kernel. We then reserve the required memory after early_init_dt_scan_memory() has identified all available physical memory. Because the two pieces of code are separated far, the readability and maintainability are reduced. So bring them together.

Suggested-by: Rob Herring <robh@kernel.org>
Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
(change the prototype of early_init_dt_check_for_usable_mem_range(), in order to use it outside)
Signed-off-by: Pingfan Liu <kernelfans@gmail.com>
Tested-by: Dave Kleikamp <dave.kleikamp@oracle.com>
Acked-by: John Donnelly <john.p.donnelly@oracle.com>

```
    Reviewed-by: Rob Herring <robh@kernel.org>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Will Deacon <will@kernel.org>
    Cc: linux-arm-kernel@lists.infradead.org
    To: devicetree@vger.kernel.org
    To: linux-efi@vger.kernel.org
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47cbc6a3aac8c720c81e4acb9628d0771988ea09
Author: William Breathitt Gray <vilhelm.gray@gmail.com>
Date:   Tue Dec 21 17:16:48 2021 +0900

    counter: 104-quad-8: Fix persistent enabled events bug

    [ Upstream commit c95cc0d95702523f8f361b802c9b7d4eeae07f5d ]

    A bug exists if the user executes a COUNTER_ADD_WATCH_IOCTL ioctl call,
    and then executes a COUNTER_DISABLE_EVENTS_IOCTL ioctl call. Disabling
    the events should disable the 104-QUAD-8 interrupts, but because of this
    bug the interrupts are not disabling.

    The reason this bug is occurring is because quad8_events_configure() is
    called when COUNTER_DISABLE_EVENTS_IOCTL is handled, but the
    next_irq_trigger[] array has not been cleared before it is checked in
    the loop.

    This patch fixes the bug by removing the next_irq_trigger array and
    instead utilizing a different algorithm of walking the events_list list
    for the current requested events. When a COUNTER_DISABLE_EVENTS_IOCTL is
    handled, events_list will be empty and thus all device channels end up
    with interrupts disabled.

    Fixes: 7aa2ba0df651 ("counter: 104-quad-8: Add IRQ support for the ACCES 104-QUAD-8")
    Cc: Syed Nayyar Waris <syednwaris@gmail.com>
    Signed-off-by: William Breathitt Gray <vilhelm.gray@gmail.com>
    Link:
https://lore.kernel.org/r/5fd5731cec1c251acee30eefb7c19160d03c9d39.1640072891.git.vilhelm.gr
ay@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07a8d2ac993ce597c3bfc576a91a320cd9f05301
Author: Trevor Wu <trevor.wu@mediatek.com>
Date:   Thu Dec 16 10:24:24 2021 +0800

    ASoC: mediatek: mt8195: correct default value

    [ Upstream commit 30e693ee82d20361f2caacca3b68c79e1a7cb16c ]

    mt8195_cg_patch is used to reset the default value of audio cg, so the
    register value could be consistent with CCF reference count.
    Nevertheless, AUDIO_TOP_CON1[1:0] is used to control an internal mux,
    and it's expected to keep the default value 0.

    This patch corrects the default value in case an unexpected behavior
    happens in the future.

    Fixes: 6746cc8582599 ("ASoC: mediatek: mt8195: add platform driver")
    Signed-off-by: Trevor Wu <trevor.wu@mediatek.com>
    Link: https://lore.kernel.org/r/20211216022424.28470-1-trevor.wu@mediatek.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15fdb770f0bbe6216745369d31d650781486310b
```

```
Author: Xiongfeng Wang <wangxiongfeng2@huawei.com>
Date:   Fri Dec 17 15:30:55 2021 +0000

    iommu/iova: Fix race between FQ timeout and teardown

    [ Upstream commit d7061627d701c90e1cac1e1e60c45292f64f3470 ]

    It turns out to be possible for hotplugging out a device to reach the
    stage of tearing down the device's group and default domain before the
    domain's flush queue has drained naturally. At this point, it is then
    possible for the timeout to expire just before the del_timer() call
    in free_iova_flush_queue(), such that we then proceed to free the FQ
    resources while fq_flush_timeout() is still accessing them on another
    CPU. Crashes due to this have been observed in the wild while removing
    NVMe devices.

    Close the race window by using del_timer_sync() to safely wait for any
    active timeout handler to finish before we start to free things. We
    already avoid any locking in free_iova_flush_queue() since the FQ is
    supposed to be inactive anyway, so the potential deadlock scenario does
    not apply.

    Fixes: 9a005a800ae8 ("iommu/iova: Add flush timer")
    Reviewed-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Xiongfeng Wang <wangxiongfeng2@huawei.com>
    [ rm: rewrite commit message ]
    Signed-off-by: Robin Murphy <robin.murphy@arm.com>
    Link:
https://lore.kernel.org/r/0a365e5b07f14b7344677ad6a9a734966a8422ce.1639753638.git.robin.murp
hy@arm.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 28f178fca05cb36500a0cf54999845c37eac4724
Author: Cezary Rojewski <cezary.rojewski@intel.com>
Date:   Thu Dec 16 12:57:39 2021 +0100

    ASoC: Intel: catpt: Test dmaengine_submit() result before moving on

    [ Upstream commit 2a9a72e290d4a4741e673f86b9fba9bfb319786d ]

    After calling dmaengine_submit(), the submitted transfer descriptor
    belongs to the DMA engine. Pointer to that descriptor may no longer be
    valid after the call and should be tested before awaiting transfer
    completion.

    Reported-by: Kevin Tian <kevin.tian@intel.com>
    Suggested-by: Dave Jiang <dave.jiang@intel.com>
    Fixes: 4fac9b31d0b9 ("ASoC: Intel: Add catpt base members")
    Signed-off-by: Cezary Rojewski <cezary.rojewski@intel.com>
    Link: https://lore.kernel.org/r/20211216115743.2130622-2-cezary.rojewski@intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 915f464d2ef24840ae8c77ae27ccca272776847d
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Tue Nov 23 18:10:38 2021 +0200

    iommu/amd: Remove useless irq affinity notifier

    [ Upstream commit 575f5cfb13c84f324f9898383fa4a5694e53c9ef ]

    iommu->intcapxt_notify field is no longer used
    after a switch to a separate domain was done
```

```
    Fixes: d1adcfbb520c ("iommu/amd: Fix IOMMU interrupt generation in X2APIC mode")
    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20211123161038.48009-6-mlevitsk@redhat.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 54a2b8e31a2bb77bd1e86021ec2bd8615a145c8e
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Tue Nov 23 18:10:37 2021 +0200

    iommu/amd: X2apic mode: mask/unmask interrupts on suspend/resume

    [ Upstream commit 1980105e3cfc2215c75b4f6b172661d675c467d1 ]

    Use IRQCHIP_MASK_ON_SUSPEND to make the core irq code to
    mask the iommu interrupt on suspend and unmask it on the resume.

    Since now the unmask function updates the INTX settings,
    that will restore them on resume from s3/s4.

    Since IRQCHIP_MASK_ON_SUSPEND is only effective for interrupts
    which are not wakeup sources, remove IRQCHIP_SKIP_SET_WAKE flag
    and instead implement a dummy .irq_set_wake which doesn't allow
    the interrupt to become a wakeup source.

    Fixes: 66929812955bb ("iommu/amd: Add support for X2APIC IOMMU interrupts")

    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20211123161038.48009-5-mlevitsk@redhat.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae648f235f8d1d4949f419f6e7a3ddedbab5e68d
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Tue Nov 23 18:10:36 2021 +0200

    iommu/amd: X2apic mode: setup the INTX registers on mask/unmask

    [ Upstream commit 4691f79d62a637958f7b5f55c232a65399500b7a ]

    This is more logically correct and will also allow us to
    to use mask/unmask logic to restore INTX setttings after
    the resume from s3/s4.

    Fixes: 66929812955bb ("iommu/amd: Add support for X2APIC IOMMU interrupts")

    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20211123161038.48009-4-mlevitsk@redhat.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7624868d67fcdd5379b8de818e198206e4a9f6e4
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Tue Nov 23 18:10:35 2021 +0200

    iommu/amd: X2apic mode: re-enable after resume

    [ Upstream commit 01b297a48a26bcb96769505ac948db4603b72bd1 ]

    Otherwise it is guaranteed to not work after the resume...

    Fixes: 66929812955bb ("iommu/amd: Add support for X2APIC IOMMU interrupts")

    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20211123161038.48009-3-mlevitsk@redhat.com
```

```
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 049e5279f61c1274e6fa9265fcb0b6689d2d0467
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Tue Nov 23 18:10:34 2021 +0200

    iommu/amd: Restore GA log/tail pointer on host resume

    [ Upstream commit a8d4a37d1bb93608501d0d0545f902061152669a ]

    This will give IOMMU GA log a chance to work after resume
    from s3/s4.

    Fixes: 8bda0cfbdc1a6 ("iommu/amd: Detect and initialize guest vAPIC log")

    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Link: https://lore.kernel.org/r/20211123161038.48009-2-mlevitsk@redhat.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c505f9c813f5d1f6080817f3b59ea6906e84faa
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Nov 22 23:21:58 2021 +0100

    dmaengine: pxa/mmp: stop referencing config->slave_id

    [ Upstream commit 134c37fa250a87a7e77c80a7c59ae16c462e46e0 ]

    The last driver referencing the slave_id on Marvell PXA and MMP platforms
    was the SPI driver, but this stopped doing so a long time ago, so the
    TODO from the earlier patch can no be removed.

    Fixes: b729bf34535e ("spi/pxa2xx: Don't use slave_id of dma_slave_config")
    Fixes: 13b3006b8ebd ("dma: mmp_pdma: add filter function")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Acked-by: Mark Brown <broonie@kernel.org>
    Link: https://lore.kernel.org/r/20211122222203.4103644-7-arnd@kernel.org
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d3280092d3dc9d2ab1e7b82b9fcb545f95122b9
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Mon Dec 13 12:16:42 2021 +0100

    mips: fix Kconfig reference to PHYS_ADDR_T_64BIT

    [ Upstream commit a670c82d9ca4f1e7385d9d6f26ff41a50fbdd944 ]

    Commit d4a451d5fc84 ("arch: remove the ARCH_PHYS_ADDR_T_64BIT config
    symbol") removes config ARCH_PHYS_ADDR_T_64BIT with all instances of that
    config refactored appropriately. Since then, it is recommended to use the
    config PHYS_ADDR_T_64BIT instead.

    Commit 171543e75272 ("MIPS: Disallow CPU_SUPPORTS_HUGEPAGES for XPA,EVA")
    introduces the expression "!(32BIT && (ARCH_PHYS_ADDR_T_64BIT || EVA))"
    for config CPU_SUPPORTS_HUGEPAGES, which unintentionally refers to the
    non-existing symbol ARCH_PHYS_ADDR_T_64BIT instead of the intended
    PHYS_ADDR_T_64BIT.

    Fix this Kconfig reference to the intended PHYS_ADDR_T_64BIT.

    This issue was identified with the script ./scripts/checkkconfigsymbols.py.
    I then reported it on the mailing list and Paul confirmed the mistake in
    the linked email thread.
```

```
    Link: https://lore.kernel.org/lkml/H8IU3R.H5QVNRA077PT@crapouillou.net/
    Suggested-by: Paul Cercueil <paul@crapouillou.net>
    Fixes: 171543e75272 ("MIPS: Disallow CPU_SUPPORTS_HUGEPAGES for XPA,EVA")
    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7d1b2966feaf215f447479c186b0e8394757020
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Mon Dec 13 12:16:35 2021 +0100

    mips: add SYS_HAS_CPU_MIPS64_R5 config for MIPS Release 5 support

    [ Upstream commit fd4eb90b164442cb1e9909f7845e12a0835ac699 ]

    Commit ab7c01fdc3cf ("mips: Add MIPS Release 5 support") adds the two
    configs CPU_MIPS32_R5 and CPU_MIPS64_R5, which depend on the corresponding
    SYS_HAS_CPU_MIPS32_R5 and SYS_HAS_CPU_MIPS64_R5, respectively.

    The config SYS_HAS_CPU_MIPS32_R5 was already introduced with commit
    c5b367835cfc ("MIPS: Add support for XPA."); the config
    SYS_HAS_CPU_MIPS64_R5, however, was never introduced.

    Hence, ./scripts/checkkconfigsymbols.py warns:

      SYS_HAS_CPU_MIPS64_R5
      Referencing files: arch/mips/Kconfig, arch/mips/include/asm/cpu-type.h

    Add the definition for config SYS_HAS_CPU_MIPS64_R5 under the assumption
    that SYS_HAS_CPU_MIPS64_R5 follows the same pattern as the existing
    SYS_HAS_CPU_MIPS32_R5 and SYS_HAS_CPU_MIPS64_R6.

    Fixes: ab7c01fdc3cf ("mips: Add MIPS Release 5 support")
    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a9b6c2f69a9ea9b1c15818fd24fca1ba3bd65d5
Author: Dillon Min <dillon.minfei@gmail.com>
Date:   Tue Oct 26 15:11:21 2021 +0800

    clk: stm32: Fix ltdc's clock turn off by clk_disable_unused() after system enter shell

    [ Upstream commit 6fc058a72f3b7b07fc4de6d66ad1f68951b00f6e ]

    stm32's clk driver register two ltdc gate clk to clk core by
    clk_hw_register_gate() and clk_hw_register_composite()

    first: 'stm32f429_gates[]', clk name is 'ltdc', which no user to use.
    second: 'stm32f429_aux_clk[]', clk name is 'lcd-tft', used by ltdc driver

    both of them point to the same offset of stm32's RCC register. after
    kernel enter console, clk core turn off ltdc's clk as 'stm32f429_gates[]'
    is no one to use. but, actually 'stm32f429_aux_clk[]' is in use.

    stm32f469/746/769 have the same issue, fix it.

    Fixes: daf2d117cbca ("clk: stm32f4: Add lcd-tft clock")
    Link: https://lore.kernel.org/linux-arm-kernel/1590564453-24499-7-git-send-email-
dillon.minfei@gmail.com/
    Link:
https://lore.kernel.org/lkml/CAPTRvHkf0cK_4ZidM17rPo99gWDmxgqFt4CDUjqFFwkOeQeFDg@mail.gmail.
com/
    Signed-off-by: Dillon Min <dillon.minfei@gmail.com>
```

```
    Reviewed-by: Patrice Chotard <patrice.chotard@foss.st.com>
    Acked-by: Gabriel Fernandez <gabriel.fernandez@foss.st.com>
    Acked-by: Stephen Boyd <sboyd@kernel.org>
    Link: https://lore.kernel.org/r/1635232282-3992-10-git-send-email-
dillon.minfei@gmail.com
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7a12cd502239ffac29faea23418eabbc0752bf3
Author: Frank Rowand <frank.rowand@sony.com>
Date:   Sun Dec 12 16:18:52 2021 -0600

    of: unittest: 64 bit dma address test requires arch support

    [ Upstream commit 9fd4cf5d3571b27d746b8ead494a3f051485b679 ]

    If an architecture does not support 64 bit dma addresses then testing
    for an expected dma address >= 0x100000000 will fail.

    Fixes: e0d072782c73 ("dma-mapping: introduce DMA range map, supplanting dma_pfn_offset")
    Signed-off-by: Frank Rowand <frank.rowand@sony.com>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20211212221852.233295-1-frowand.list@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b29fbc2a610e953eb9ec4ebbf513514377e5f83e
Author: Jim Quinlan <jim2101024@gmail.com>
Date:   Fri Dec 10 13:46:35 2021 -0500

    of: unittest: fix warning on PowerPC frame size warning

    [ Upstream commit a8d61a9112ad0c9216ab45d050991e07bc4f3408 ]

    The struct device variable "dev_bogus" was triggering this warning
    on a PowerPC build:

        drivers/of/unittest.c: In function 'of_unittest_dma_ranges_one.constprop':
        [...] >> The frame size of 1424 bytes is larger than 1024 bytes
                 [-Wframe-larger-than=]

    This variable is now dynamically allocated.

    Fixes: e0d072782c734 ("dma-mapping: introduce DMA range map, supplanting
dma_pfn_offset")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Jim Quinlan <jim2101024@gmail.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Frank Rowand <frank.rowand@sony.com>
    Reviewed-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20211210184636.7273-2-jim2101024@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b015e4dca7a11bcd49bf8cf82e11dc7ea0d76cd0
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Dec 15 11:15:50 2021 +0800

    ASoC: rt5663: Handle device_property_read_u32_array error codes

    [ Upstream commit 2167c0b205960607fb136b4bb3c556a62be1569a ]

    The return value of device_property_read_u32_array() is not always 0.
    To catch the exception in case that devm_kzalloc failed and the
    rt5663->imp_table was NULL, which caused the failure of
    device_property_read_u32_array.
```

```
    Fixes: 450f0f6a8fb4 ("ASoC: rt5663: Add the manual offset field to compensate the DC
offset")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Link: https://lore.kernel.org/r/20211215031550.70702-1-jiasheng@iscas.ac.cn
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 34ca47bed842dfd97f293118b1bc71e31de43869
Author: Avihai Horon <avihaih@nvidia.com>
Date:   Thu Dec 9 15:16:07 2021 +0200

    RDMA/cma: Let cma_resolve_ib_dev() continue search even after empty entry

    [ Upstream commit 20679094a0161c94faf77e373fa3f7428a8e14bd ]

    Currently, when cma_resolve_ib_dev() searches for a matching GID it will
    stop searching after encountering the first empty GID table entry. This
    behavior is wrong since neither IB nor RoCE spec enforce tightly packed
    GID tables.

    For example, when the matching valid GID entry exists at index N, and if a
    GID entry is empty at index N-1, cma_resolve_ib_dev() will fail to find
    the matching valid entry.

    Fix it by making cma_resolve_ib_dev() continue searching even after
    encountering missing entries.

    Fixes: f17df3b0dede ("RDMA/cma: Add support for AF_IB to rdma_resolve_addr()")
    Link:
https://lore.kernel.org/r/b7346307e3bb396c43d67d924348c6c496493991.1639055490.git.leonro@nvi
dia.com
    Signed-off-by: Avihai Horon <avihaih@nvidia.com>
    Reviewed-by: Mark Zhang <markzhang@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a72907c4ed2e04dd7c7821b31b68963fdaaa865
Author: Avihai Horon <avihaih@nvidia.com>
Date:   Thu Dec 9 15:16:06 2021 +0200

    RDMA/core: Let ib_find_gid() continue search even after empty entry

    [ Upstream commit 483d805191a23191f8294bbf9b4e94836f5d92e4 ]

    Currently, ib_find_gid() will stop searching after encountering the first
    empty GID table entry. This behavior is wrong since neither IB nor RoCE
    spec enforce tightly packed GID tables.

    For example, when a valid GID entry exists at index N, and if a GID entry
    is empty at index N-1, ib_find_gid() will fail to find the valid entry.

    Fix it by making ib_find_gid() continue searching even after encountering
    missing entries.

    Fixes: 5eb620c81ce3 ("IB/core: Add helpers for uncached GID and P_Key searches")
    Link:
https://lore.kernel.org/r/e55d331b96cecfc2cf19803d16e7109ea966882d.1639055490.git.leonro@nvi
dia.com
    Signed-off-by: Avihai Horon <avihaih@nvidia.com>
    Reviewed-by: Mark Zhang <markzhang@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 4dfb268a9f6fb0c54e8aa4f5921756afed016f95
Author: Rob Clark <robdclark@chromium.org>
Date:   Mon Nov 8 09:17:23 2021 -0800

    iommu/arm-smmu-qcom: Fix TTBR0 read

    [ Upstream commit c31112fbd4077a51a14ff338038c82e9571dc821 ]

    It is a 64b register, lets not lose the upper bits.

    Fixes: ab5df7b953d8 ("iommu/arm-smmu-qcom: Add an adreno-smmu-priv callback to get
pagefault info")
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211108171724.470973-1-robdclark@gmail.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2bfd5c28a8f0730a29eed19f89c6631200bad782
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Wed Dec 8 17:36:52 2021 +0000

    powerpc/powermac: Add additional missing lockdep_register_key()

    [ Upstream commit b149d5d45ac9171ed699a256f026c8ebef901112 ]

    Commit df1f679d19ed ("powerpc/powermac: Add missing
    lockdep_register_key()") fixed a problem that was causing a WARNING.

    There are two other places in the same file with the same problem
    originating from commit 9e607f72748d ("i2c_powermac: shut up lockdep
    warning").

    Add missing lockdep_register_key()

    Fixes: 9e607f72748d ("i2c_powermac: shut up lockdep warning")
    Reported-by: Erhard Furtner <erhard_f@mailbox.org>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Depends-on: df1f679d19ed ("powerpc/powermac: Add missing lockdep_register_key()")
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=200055
    Link:
https://lore.kernel.org/r/2c7e421874e21b2fb87813d768cf662f630c2ad4.1638984999.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4dd0a3184b5cfd3926d39cefa77e0611e5a548f4
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Mon Dec 6 23:27:26 2021 +0100

    PCI/MSI: Fix pci_irq_vector()/pci_irq_get_affinity()

    [ Upstream commit 29bbc35e29d9b6347780dcacde2deb4b39344167 ]

    pci_irq_vector() and pci_irq_get_affinity() use the list position to find the
    MSI-X descriptor at a given index. That's correct for the normal case where
    the entry number is the same as the list position.

    But it's wrong for cases where MSI-X was allocated with an entries array
    describing sparse entry numbers into the hardware message descriptor
    table. That's inconsistent at best.

    Make it always check the entry number because that's what the zero base
    index really means. This change won't break existing users which use a
```

```
    sparse entries array for allocation because these users retrieve the Linux
    interrupt number from the entries array after allocation and none of them
    uses pci_irq_vector() or pci_irq_get_affinity().

    Fixes: aff171641d18 ("PCI: Provide sensible IRQ vector alloc/free routines")
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Juergen Gross <jgross@suse.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Acked-by: Bjorn Helgaas <bhelgaas@google.com>
    Link: https://lore.kernel.org/r/20211206210223.929792157@linutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 10862d6e49ba85f34e1536f8427d6a5f2db41f5a
Author: Kamal Heib <kamalheib1@gmail.com>
Date:   Mon Dec 6 22:13:14 2021 +0200

    RDMA/qedr: Fix reporting max_{send/recv}_wr attrs

    [ Upstream commit b1a4da64bfc189510e08df1ccb1c589e667dc7a3 ]

    Fix the wrongly reported max_send_wr and max_recv_wr attributes for user
    QP by making sure to save their valuse on QP creation, so when query QP is
    called the attributes will be reported correctly.

    Fixes: cecbcddf6461 ("qedr: Add support for QP verbs")
    Link: https://lore.kernel.org/r/20211206201314.124947-1-kamalheib1@gmail.com
    Signed-off-by: Kamal Heib <kamalheib1@gmail.com>
    Acked-by: Michal Kalderon <michal.kalderon@marvell.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5eae90bbe305f3bee931dd9abb5a89927ece3764
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Fri Dec 3 15:19:39 2021 -0800

    scsi: ufs: Fix race conditions related to driver data

    [ Upstream commit 21ad0e49085deb22c094f91f9da57319a97188e4 ]

    The driver data pointer must be set before any callbacks are registered
    that use that pointer. Hence move the initialization of that pointer from
    after the ufshcd_init() call to inside ufshcd_init().

    Link: https://lore.kernel.org/r/20211203231950.193369-7-bvanassche@acm.org
    Fixes: 3b1d05807a9a ("[SCSI] ufs: Segregate PCI Specific Code")
    Reported-by: Alexey Dobriyan <adobriyan@gmail.com>
    Tested-by: Bean Huo <beanhuo@micron.com>
    Reviewed-by: Bean Huo <beanhuo@micron.com>
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c0645486821c4d97e2ce45473924fe821c18a52e
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Fri Dec 3 15:19:34 2021 -0800

    scsi: core: Fix scsi_device_max_queue_depth()

    [ Upstream commit 4bc3bffc1a885eb5cb259e4a25146a4c7b1034e3 ]

    The comment above scsi_device_max_queue_depth() and also the description of
    commit ca4453213951 ("scsi: core: Make sure sdev->queue_depth is <=
    max(shost->can_queue, 1024)") contradict the implementation of the function
    scsi_device_max_queue_depth(). Additionally, the maximum queue depth of a
    SCSI LUN never exceeds host->can_queue. Fix scsi_device_max_queue_depth()
```

```
    by changing max_t() into min_t().

    Link: https://lore.kernel.org/r/20211203231950.193369-2-bvanassche@acm.org
    Fixes: ca4453213951 ("scsi: core: Make sure sdev->queue_depth is <= max(shost-
>can_queue, 1024)")
    Cc: Hannes Reinecke <hare@suse.de>
    Cc: Sumanesh Samanta <sumanesh.samanta@broadcom.com>
    Tested-by: Bean Huo <beanhuo@micron.com>
    Reviewed-by: Ming Lei <ming.lei@redhat.com>
    Reviewed-by: Bean Huo <beanhuo@micron.com>
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2cc5e97a20a077b531ee098a7dde0e9c2f610b7e
Author: Kai Vehmanen <kai.vehmanen@linux.intel.com>
Date:   Fri Dec 3 17:47:21 2021 +0200

    ASoC: SOF: Intel: fix build issue related to CODEC_PROBE_ENTRIES

    [ Upstream commit 9a83dfcc5ae8230fbf12b63e281d5bb8450ec0e7 ]

    Fix following error:
    sound/soc/sof/intel/hda-codec.c:132:35: error: use of undeclared identifier
'CODEC_PROBE_RETRIES'

    Found with config: i386-randconfig-r033-20211202
    (https://download.01.org/0day-ci/archive/20211203/202112031943.Twg19fWT-
lkp@intel.com/config)

    Fixes: 046aede2f847 ("ASoC: SOF: Intel: Retry codec probing if it fails")
    Reported-by: kernel test robot <lkp@intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20211203154721.923496-1-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc31a28f342e1383f4a93b0379dd7cc03a7c6b87
Author: Hector Martin <marcan@marcan.st>
Date:   Sat Nov 20 12:13:43 2021 +0900

    iommu/io-pgtable-arm: Fix table descriptor paddr formatting

    [ Upstream commit 9abe2ac834851a7d0b0756e295cf7a292c45ca53 ]

    Table descriptors were being installed without properly formatting the
    address using paddr_to_iopte, which does not match up with the
    iopte_deref in __arm_lpae_map. This is incorrect for the LPAE pte
    format, as it does not handle the high bits properly.

    This was found on Apple T6000 DARTs, which require a new pte format
    (different shift); adding support for that to
    paddr_to_iopte/iopte_to_paddr caused it to break badly, as even <48-bit
    addresses would end up incorrect in that case.

    Fixes: 6c89928ff7a0 ("iommu/io-pgtable-arm: Support 52-bit physical address")
    Acked-by: Robin Murphy <robin.murphy@arm.com>
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Link: https://lore.kernel.org/r/20211120031343.88034-1-marcan@marcan.st
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a21c7a03a43c10c64f80b36893964b1d3036f1cb
Author: Lu Baolu <baolu.lu@linux.intel.com>
```

```
Date:   Mon Nov 8 14:13:49 2021 +0800

    iommu: Extend mutex lock scope in iommu_probe_device()

    [ Upstream commit 556f99ac886635e8da15528995f06d1d7028cfca ]

    Extend the scope of holding group->mutex so that it can cover the default
    domain check/attachment and direct mappings of reserved regions.

    Cc: Ashish Mhetre <amhetre@nvidia.com>
    Fixes: 211ff31b3d33b ("iommu: Fix race condition during default domain allocation")
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Link: https://lore.kernel.org/r/20211108061349.1985579-1-baolu.lu@linux.intel.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1de65f0017f69f9f9a5d9aa13778a445623ccfa5
Author: Stafford Horne <shorne@gmail.com>
Date:   Sat Dec 4 07:10:18 2021 +0900

    openrisc: Add clone3 ABI wrapper

    [ Upstream commit 433fe39f674d58bc7a3e8254a5d2ffc290b7e04e ]

    Like fork and clone the clone3 syscall needs a wrapper to save callee
    saved registers, which is required by the OpenRISC ABI.  This came up
    after auditing code following a discussion with Rob Landley and Arnd
    Bergmann [0].

    Tested with the clone3 kselftests and there were no issues.

    [0] https://lore.kernel.org/all/41206fc7-f8ce-98aa-3718-
ba3e1431e320@landley.net/T/#m9c0cdb2703813b9df4da04cf6b30de1f1aa89944

    Fixes: 07e83dfbe16c ("openrisc: Enable the clone3 syscall")
    Cc: Rob Landley <rob@landley.net>
    Cc: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Stafford Horne <shorne@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66e12f5b3a9733f941893a00753b10498724607d
Author: Todd Kjos <tkjos@google.com>
Date:   Tue Nov 30 10:51:50 2021 -0800

    binder: avoid potential data leakage when copying txn

    [ Upstream commit 6d98eb95b450a75adb4516a1d33652dc78d2b20c ]

    Transactions are copied from the sender to the target
    first and objects like BINDER_TYPE_PTR and BINDER_TYPE_FDA
    are then fixed up. This means there is a short period where
    the sender's version of these objects are visible to the
    target prior to the fixups.

    Instead of copying all of the data first, copy data only
    after any needed fixups have been applied.

    Fixes: 457b9a6f09f0 ("Staging: android: add binder driver")
    Reviewed-by: Martijn Coenen <maco@android.com>
    Acked-by: Christian Brauner <christian.brauner@ubuntu.com>
    Signed-off-by: Todd Kjos <tkjos@google.com>
    Link: https://lore.kernel.org/r/20211130185152.437403-3-tkjos@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 19066af3763c836782728539adaee93ad96759df
Author: Todd Kjos <tkjos@google.com>
Date:   Tue Nov 30 10:51:49 2021 -0800

    binder: fix handling of error during copy

    [ Upstream commit fe6b1869243f23a485a106c214bcfdc7aa0ed593 ]

    If a memory copy function fails to copy the whole buffer,
    a positive integar with the remaining bytes is returned.
    In binder_translate_fd_array() this can result in an fd being
    skipped due to the failed copy, but the loop continues
    processing fds since the early return condition expects a
    negative integer on error.

    Fix by returning "ret > 0 ? -EINVAL : ret" to handle this case.

    Fixes: bb4a2e48d510 ("binder: return errors from buffer copy functions")
    Suggested-by: Dan Carpenter <dan.carpenter@oracle.com>
    Acked-by: Christian Brauner <christian.brauner@ubuntu.com>
    Signed-off-by: Todd Kjos <tkjos@google.com>
    Link: https://lore.kernel.org/r/20211130185152.437403-2-tkjos@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fef06f38906a04c3bd71756ff5cf6bf5ba288dc
Author: Kees Cook <keescook@chromium.org>
Date:   Fri Dec 3 00:42:06 2021 -0800

    char/mwave: Adjust io port register size

    [ Upstream commit f5912cc19acd7c24b2dbf65a6340bf194244f085 ]

    Using MKWORD() on a byte-sized variable results in OOB read. Expand the
    size of the reserved area so both MKWORD and MKBYTE continue to work
    without overflow. Silences this warning on a -Warray-bounds build:

    drivers/char/mwave/3780i.h:346:22: error: array subscript 'short unsigned int[0]' is
partly outside array bounds of 'DSP_ISA_SLAVE_CONTROL[1]' [-Werror=array-bounds]
      346 | #define MKWORD(var) (*((unsigned short *)(&var)))
          |                      ~^~~~~~~~~~~~~~~~~~~~~~~~~~
    drivers/char/mwave/3780i.h:356:40: note: in definition of macro 'OutWordDsp'
      356 | #define OutWordDsp(index,value)   outw(value,usDspBaseIO+index)
          |                                        ^~~~~
    drivers/char/mwave/3780i.c:373:41: note: in expansion of macro 'MKWORD'
      373 |         OutWordDsp(DSP_IsaSlaveControl, MKWORD(rSlaveControl));
          |                                         ^~~~~~
    drivers/char/mwave/3780i.c:358:31: note: while referencing 'rSlaveControl'
      358 |         DSP_ISA_SLAVE_CONTROL rSlaveControl;
          |                               ^~~~~~~~~~~~

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20211203084206.3104326-1-keescook@chromium.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b25a705cea16f5849618d9475913d1472d0b6157
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Thu Nov 25 23:27:27 2021 +0200

    misc: at25: Make driver OF independent again

    [ Upstream commit 5b557298d7d09cce04e0565a535fbca63661724a ]
```

```
    The commit f60e7074902a ("misc: at25: Make use of device property API")
    made a good job by enabling the driver for non-OF platforms, but the
    recent commit 604288bc6196 ("nvmem: eeprom: at25: fix type compiler warnings")
    brought that back.

    Restore greatness of the driver once again.

    Fixes: eab61fb1cc2e ("nvmem: eeprom: at25: fram discovery simplification")
    Fixes: fd307a4ad332 ("nvmem: prepare basics for FRAM support")
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20211125212729.86585-2-andriy.shevchenko@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d0607af54999ef38fbafd315d4202042c8ea0f7d
Author: Lucas Tanure <tanureal@opensource.cirrus.com>
Date:   Wed Dec 1 18:00:03 2021 +0000

    ASoC: amd: Fix dependency for SPI master

    [ Upstream commit 19a628d8f1a6c16263d8037a918427207c8a95a0 ]

    Set SPI_MASTER as dependency as is using CS35L41 SPI driver

    Fixes: 96792fdd77cd1 ("ASoC: amd: enable vangogh platform machine driver build")

    Signed-off-by: Lucas Tanure <tanureal@opensource.cirrus.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Link: https://lore.kernel.org/r/20211201180004.1402156-1-tanureal@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e5cfbde83d9e7e1bccdad28d0f65a5fa78fb7e5f
Author: Takashi Iwai <tiwai@suse.de>
Date:   Thu Dec 2 09:38:33 2021 +0100

    ALSA: usb-audio: Drop superfluous '0' in Presonus Studio 1810c's ID

    [ Upstream commit 1e583aef12aa74afd37c1418255cc4b74e023236 ]

    The vendor ID of Presonus Studio 1810c had a superfluous '0' in its
    USB ID.  Drop it.

    Fixes: 8dc5efe3d17c ("ALSA: usb-audio: Add support for Presonus Studio 1810c")
    Link: https://lore.kernel.org/r/20211202083833.17784-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a9fdca41ada61af0a0de121e16fa060727471db
Author: Bixuan Cui <cuibixuan@linux.alibaba.com>
Date:   Wed Dec 1 16:58:54 2021 +0800

    ALSA: oss: fix compile error when OSS_DEBUG is enabled

    [ Upstream commit 8e7daf318d97f25e18b2fc7eb5909e34cd903575 ]

    Fix compile error when OSS_DEBUG is enabled:
        sound/core/oss/pcm_oss.c: In function 'snd_pcm_oss_set_trigger':
        sound/core/oss/pcm_oss.c:2055:10: error: 'substream' undeclared (first
        use in this function); did you mean 'csubstream'?
          pcm_dbg(substream->pcm, "pcm_oss: trigger = 0x%x\n", trigger);
                  ^


    Fixes: 61efcee8608c ("ALSA: oss: Use standard printk helpers")
```

```
    Signed-off-by: Bixuan Cui <cuibixuan@linux.alibaba.com>
    Link: https://lore.kernel.org/r/1638349134-110369-1-git-send-email-
cuibixuan@linux.alibaba.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b7a7602b2e041f156cf7a1b4fb8cbe0eb506dc1
Author: Waiman Long <longman@redhat.com>
Date:   Thu Nov 18 14:14:36 2021 -0500

    clocksource: Avoid accidental unstable marking of clocksources

    [ Upstream commit c86ff8c55b8ae68837b2fa59dc0c203907e9a15f ]

    Since commit db3a34e17433 ("clocksource: Retry clock read if long delays
    detected") and commit 2e27e793e280 ("clocksource: Reduce clocksource-skew
    threshold"), it is found that tsc clocksource fallback to hpet can
    sometimes happen on both Intel and AMD systems especially when they are
    running stressful benchmarking workloads. Of the 23 systems tested with
    a v5.14 kernel, 10 of them have switched to hpet clock source during
    the test run.

    The result of falling back to hpet is a drastic reduction of performance
    when running benchmarks. For example, the fio performance tests can
    drop up to 70% whereas the iperf3 performance can drop up to 80%.

    4 hpet fallbacks happened during bootup. They were:

      [    8.749399] clocksource: timekeeping watchdog on CPU13: hpet read-back delay of
263750ns, attempt 4, marking unstable
      [   12.044610] clocksource: timekeeping watchdog on CPU19: hpet read-back delay of
186166ns, attempt 4, marking unstable
      [   17.336941] clocksource: timekeeping watchdog on CPU28: hpet read-back delay of
182291ns, attempt 4, marking unstable
      [   17.518565] clocksource: timekeeping watchdog on CPU34: hpet read-back delay of
252196ns, attempt 4, marking unstable

    Other fallbacks happen when the systems were running stressful
    benchmarks. For example:

      [ 2685.867873] clocksource: timekeeping watchdog on CPU117: hpet read-back delay of
57269ns, attempt 4, marking unstable
      [46215.471228] clocksource: timekeeping watchdog on CPU8: hpet read-back delay of
61460ns, attempt 4, marking unstable

    Commit 2e27e793e280 ("clocksource: Reduce clocksource-skew threshold"),
    changed the skew margin from 100us to 50us. I think this is too small
    and can easily be exceeded when running some stressful workloads on a
    thermally stressed system.  So it is switched back to 100us.

    Even a maximum skew margin of 100us may be too small in for some systems
    when booting up especially if those systems are under thermal stress. To
    eliminate the case that the large skew is due to the system being too
    busy slowing down the reading of both the watchdog and the clocksource,
    an extra consecutive read of watchdog clock is being done to check this.

    The consecutive watchdog read delay is compared against
    WATCHDOG_MAX_SKEW/2. If the delay exceeds the limit, we assume that
    the system is just too busy. A warning will be printed to the console
    and the clock skew check is skipped for this round.

    Fixes: db3a34e17433 ("clocksource: Retry clock read if long delays detected")
    Fixes: 2e27e793e280 ("clocksource: Reduce clocksource-skew threshold")
    Signed-off-by: Waiman Long <longman@redhat.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
```

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

commit a1ae41f691c7738658a3c87df5c7ad44381a75c1
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Tue Nov 30 09:42:37 2021 +0100

    powerpc/32s: Fix shift-out-of-bounds in KASAN init

    [ Upstream commit af11dee4361b3519981fa04d014873f9d9edd6ac ]

    ================================================================================
    UBSAN: shift-out-of-bounds in arch/powerpc/mm/kasan/book3s_32.c:22:23
    shift exponent -1 is negative
    CPU: 0 PID: 0 Comm: swapper Not tainted 5.15.5-gentoo-PowerMacG4 #9
    Call Trace:
    [c214be60] [c0ba0048] dump_stack_lvl+0x80/0xb0 (unreliable)
    [c214be80] [c0b99288] ubsan_epilogue+0x10/0x5c
    [c214be90] [c0b98fe0] __ubsan_handle_shift_out_of_bounds+0x94/0x138
    [c214bf00] [c1c0f010] kasan_init_region+0xd8/0x26c
    [c214bf30] [c1c0ed84] kasan_init+0xc0/0x198
    [c214bf70] [c1c08024] setup_arch+0x18/0x54c
    [c214bfc0] [c1c037f0] start_kernel+0x90/0x33c
    [c214bff0] [00003610] 0x3610
    ================================================================================

    This happens when the directly mapped memory is a power of 2.

    Fix it by checking the shift and set the result to 0 when shift is -1

    Fixes: 7974c4732642 ("powerpc/32s: Implement dedicated kasan_init_region()")
    Reported-by: Erhard Furtner <erhard_f@mailbox.org>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=215169
    Link:
https://lore.kernel.org/r/15cbc3439d4ad988b225e2119ec99502a5cc6ad3.1638261744.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c31544cded7d87b7b9f0c6df601eb1af8ea04fa
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Tue Nov 30 11:10:43 2021 +0100

    powerpc/modules: Don't WARN on first module allocation attempt

    [ Upstream commit f1797e4de1146009c888bcf8b6bb6648d55394f1 ]

    module_alloc() first tries to allocate module text within 24 bits direct
    jump from kernel text, and tries a wider allocation if first one fails.

    When first allocation fails the following is observed in kernel logs:

      vmap allocation for size 2400256 failed: use vmalloc=<size> to increase size
      systemd-udevd: vmalloc error: size 2395133, vm_struct allocation failed,
mode:0xcc0(GFP_KERNEL), nodemask=(null)
      CPU: 0 PID: 127 Comm: systemd-udevd Tainted: G        W         5.15.5-gentoo-
PowerMacG4 #9
      Call Trace:
      [e2a53a50] [c0ba0048] dump_stack_lvl+0x80/0xb0 (unreliable)
      [e2a53a70] [c0540128] warn_alloc+0x11c/0x2b4
      [e2a53b50] [c0531be8] __vmalloc_node_range+0xd8/0x64c
      [e2a53c10] [c00338c0] module_alloc+0xa0/0xac
      [e2a53c40] [c027a368] load_module+0x2ae0/0x8148
      [e2a53e30] [c027fc78] sys_finit_module+0xfc/0x130
      [e2a53f30] [c0035098] ret_from_syscall+0x0/0x28
```

```
      ...

    Add __GFP_NOWARN flag to first allocation so that no warning appears
    when it fails.

    Reported-by: Erhard Furtner <erhard_f@mailbox.org>
    Fixes: 2ec13df16704 ("powerpc/modules: Load modules closer to kernel text")
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/93c9b84d6ec76aaf7b4f03468e22433a6d308674.1638267035.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 215a90ce3754fe509efbce6b73a4bb643c7e7528
Author: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
Date:   Wed Jul 21 01:48:29 2021 -0400

    powerpc/perf: Fix PMU callbacks to clear pending PMI before resetting an overflown PMC

    [ Upstream commit 2c9ac51b850d84ee496b0a5d832ce66d411ae552 ]

    Running perf fuzzer showed below in dmesg logs:
      "Can't find PMC that caused IRQ"

    This means a PMU exception happened, but none of the PMC's (Performance
    Monitor Counter) were found to be overflown. There are some corner cases
    that clears the PMCs after PMI gets masked. In such cases, the perf
    interrupt handler will not find the active PMC values that had caused
    the overflow and thus leads to this message while replaying.

    Case 1: PMU Interrupt happens during replay of other interrupts and
    counter values gets cleared by PMU callbacks before replay:

    During replay of interrupts like timer, __do_irq() and doorbell
    exception, we conditionally enable interrupts via may_hard_irq_enable().
    This could potentially create a window to generate a PMI. Since irq soft
    mask is set to ALL_DISABLED, the PMI will get masked here. We could get
    IPIs run before perf interrupt is replayed and the PMU events could
    be deleted or stopped. This will change the PMU SPR values and resets
    the counters. Snippet of ftrace log showing PMU callbacks invoked in
    __do_irq():

      <idle>-0 [051] dns. 132025441306354: __do_irq <-call_do_irq
      <idle>-0 [051] dns. 132025441306430: irq_enter <-__do_irq
      <idle>-0 [051] dns. 132025441306503: irq_enter_rcu <-__do_irq
      <idle>-0 [051] dnH. 132025441306599: xive_get_irq <-__do_irq
      <<>>
      <idle>-0 [051] dnH. 132025441307770: generic_smp_call_function_single_interrupt <-
smp_ipi_demux_relaxed
      <idle>-0 [051] dnH. 132025441307839: flush_smp_call_function_queue <-
smp_ipi_demux_relaxed
      <idle>-0 [051] dnH. 132025441308057: _raw_spin_lock <-event_function
      <idle>-0 [051] dnH. 132025441308206: power_pmu_disable <-perf_pmu_disable
      <idle>-0 [051] dnH. 132025441308337: power_pmu_del <-event_sched_out
      <idle>-0 [051] dnH. 132025441308407: power_pmu_read <-power_pmu_del
      <idle>-0 [051] dnH. 132025441308477: read_pmc <-power_pmu_read
      <idle>-0 [051] dnH. 132025441308590: isa207_disable_pmc <-power_pmu_del
      <idle>-0 [051] dnH. 132025441308663: write_pmc <-power_pmu_del
      <idle>-0 [051] dnH. 132025441308787: power_pmu_event_idx <-perf_event_update_userpage
      <idle>-0 [051] dnH. 132025441308859: rcu_read_unlock_strict <-
perf_event_update_userpage
      <idle>-0 [051] dnH. 132025441308975: power_pmu_enable <-perf_pmu_enable
      <<>>
      <idle>-0 [051] dnH. 132025441311108: irq_exit <-__do_irq
```

```
       <idle>-0 [051] dns. 132025441311319: performance_monitor_exception <-
  replay_soft_interrupts

      Case 2: PMI's masked during local_* operations, example local_add(). If
      the local_add() operation happens within a local_irq_save(), replay of
      PMI will be during local_irq_restore(). Similar to case 1, this could
      also create a window before replay where PMU events gets deleted or
      stopped.

      Fix it by updating the PMU callback function power_pmu_disable() to
      check for pending perf interrupt. If there is an overflown PMC and
      pending perf interrupt indicated in paca, clear the PMI bit in paca to
      drop that sample. Clearing of PMI bit is done in power_pmu_disable()
      since disable is invoked before any event gets deleted/stopped. With
      this fix, if there are more than one event running in the PMU, there is
      a chance that we clear the PMI bit for the event which is not getting
      deleted/stopped. The other events may still remain active. Hence to make
      sure we don't drop valid sample in such cases, another check is added in
      power_pmu_enable. This checks if there is an overflown PMC found among
      the active events and if so enable back the PMI bit. Two new helper
      functions are introduced to clear/set the PMI, ie
      clear_pmi_irq_pending() and set_pmi_irq_pending(). Helper function
      pmi_irq_pending() is introduced to give a warning if there is pending
      PMI bit in paca, but no PMC is overflown.

      Also there are corner cases which result in performance monitor
      interrupts being triggered during power_pmu_disable(). This happens
      since PMXE bit is not cleared along with disabling of other MMCR0 bits
      in the pmu_disable. Such PMI's could leave the PMU running and could
      trigger PMI again which will set MMCR0 PMAO bit. This could lead to
      spurious interrupts in some corner cases. Example, a timer after
      power_pmu_del() which will re-enable interrupts and triggers a PMI again
      since PMAO bit is still set. But fails to find valid overflow since PMC
      was cleared in power_pmu_del(). Fix that by disabling PMXE along with
      disabling of other MMCR0 bits in power_pmu_disable().

      We can't just replay PMI any time. Hence this approach is preferred
      rather than replaying PMI before resetting overflown PMC. Patch also
      documents core-book3s on a race condition which can trigger these PMC
      messages during idle path in PowerNV.

      Fixes: f442d004806e ("powerpc/64s: Add support to mask perf interrupts and replay them")
      Reported-by: Nageswara R Sastry <nasastry@in.ibm.com>
      Suggested-by: Nicholas Piggin <npiggin@gmail.com>
      Suggested-by: Madhavan Srinivasan <maddy@linux.ibm.com>
      Signed-off-by: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
      Tested-by: Nageswara R Sastry <rnsastry@linux.ibm.com>
      Reviewed-by: Nicholas Piggin <npiggin@gmail.com>
      [mpe: Make pmi_irq_pending() return bool, reflow/reword some comments]
      Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
      Link: https://lore.kernel.org/r/1626846509-1350-2-git-send-email-
  atrajeev@linux.vnet.ibm.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 33bbf23379952cf032352887470e25e546bc6f17
  Author: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
  Date:   Tue Nov 9 11:30:45 2021 +0100

      dt-bindings: thermal: Fix definition of cooling-maps contribution property

      [ Upstream commit 49bcb1506f2e095262c01bda7fd1c0db524c91e2 ]

      When converting the thermal-zones bindings to yaml the definition of the
      contribution property changed. The intention is the same, an integer
      value expressing a ratio of a sum on how much cooling is provided by the
```

```
    device to the zone. But after the conversion the integer value is
    limited to the range 0 to 100 and expressed as a percentage.

    This is problematic for two reasons.

    - This do not match how the binding is used. Out of the 18 files that
      make use of the property only two (ste-dbx5x0.dtsi and
      ste-hrefv60plus.dtsi) sets it at a value that satisfy the binding,
      100. The remaining 16 files set the value higher and fail to validate.

    - Expressing the value as a percentage instead of a ratio of the sum is
      confusing as there is nothing to enforce the sum in the zone is not
      greater then 100.

    This patch restore the pre yaml conversion description and removes the
    value limitation allowing the usage of the bindings to validate.

    Fixes: 1202a442a31fd2e5 ("dt-bindings: thermal: Add yaml bindings for thermal zones")
    Reported-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Signed-off-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Link: https://lore.kernel.org/r/20211109103045.1403686-1-
niklas.soderlund+renesas@ragnatech.se
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a96b8a1ccc59e14c57dc6c5bd8b5be560fc48622
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Nov 24 23:40:01 2021 +0100

    ALSA: hda: Make proper use of timecounter

    [ Upstream commit 6dd21ad81bf96478db3403b1bbe251c0612d0431 ]

    HDA uses a timecounter to read a hardware clock running at 24 MHz. The
    conversion factor is set with a mult value of 125 and a shift value of 0,
    which is not converting the hardware clock to nanoseconds, it is converting
    to 1/3 nanoseconds because the conversion factor from 24Mhz to nanoseconds
    is 125/3. The usage sites divide the "nanoseconds" value returned by
    timecounter_read() by 3 to get a real nanoseconds value.

    There is a lengthy comment in azx_timecounter_init() explaining this
    choice. That comment makes blatantly wrong assumptions about how
    timecounters work and what can overflow.

    The comment says:

         * Applying the 1/3 factor as part of the multiplication
         * requires at least 20 bits for a decent precision, however
         * overflows occur after about 4 hours or less, not a option.

    timecounters operate on time deltas between two readouts of a clock and use
    the mult/shift pair to calculate a precise nanoseconds value:

        delta_nsec = (delta_clock * mult) >> shift;

    The fractional part is also taken into account and preserved to prevent
    accumulated rounding errors. For details see cyclecounter_cyc2ns().

    The mult/shift pair has to be chosen so that the multiplication of the
    maximum expected delta value does not result in a 64bit overflow. As the
    counter wraps around on 32bit, the maximum observable delta between two
    reads is (1 << 32) - 1 which is about 178.9 seconds.

    That in turn means the maximum multiplication factor which fits into an u32
    will not cause a 64bit overflow ever because it's guaranteed that:
```

```
          ((1 << 32) - 1) ^ 2 < (1 << 64)

      The resulting correct multiplication factor is 2796202667 and the shift
      value is 26, i.e. 26 bit precision. The overflow of the multiplication
      would happen exactly at a clock readout delta of 6597069765 which is way
      after the wrap around of the hardware clock at around 274.8 seconds which
      is off from the claimed 4 hours by more than an order of magnitude.

      If the counter ever wraps around the last read value then the calculation
      is off by the number of wrap arounds times 178.9 seconds because the
      overflow cannot be observed.

      Use clocks_calc_mult_shift(), which calculates the most accurate mult/shift
      pair based on the given clock frequency, and remove the bogus comment along
      with the divisions at the readout sites.

      Fixes: 5d890f591d15 ("ALSA: hda: support for wallclock timestamps")
      Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
      Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
      Link: https://lore.kernel.org/r/871r35kwji.ffs@tglx
      Signed-off-by: Takashi Iwai <tiwai@suse.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fee49d76691c51437d0b14da39b8d9cba8bc0cd1
  Author: Jack Wang <jinpu.wang@ionos.com>
  Date:   Wed Nov 24 09:10:40 2021 +0100

      RDMA/rtrs-clt: Fix the initial value of min_latency

      [ Upstream commit 925cac6358677d3d64f9b25f205eeb3d31c9f7f8 ]

      The type of min_latency is ktime_t, so use KTIME_MAX to initialize the
      initial value.

      Fixes: dc3b66a0ce70 ("RDMA/rtrs-clt: Add a minimum latency multipath policy")
      Link: https://lore.kernel.org/r/20211124081040.19533-1-jinpu.wang@ionos.com
      Signed-off-by: Jack Wang <jinpu.wang@ionos.com>
      Reviewed-by: Guoqing Jiang <Guoqing.Jiang@linux.dev>
      Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit f29140d9dbe2a72d8dadf891218547d2ffcfa1d7
  Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
  Date:   Thu Nov 25 10:51:58 2021 +0100

      ASoC: codecs: wcd938x: add SND_SOC_WCD938_SDW to codec list instead

      [ Upstream commit 2039cc1da4bee1fd0df644e26b28ed769cd32a81 ]

      Commit 045442228868 ("ASoC: codecs: wcd938x: add audio routing and
      Kconfig") adds SND_SOC_WCD937X, which does not exist, and
      SND_SOC_WCD938X, which seems not really to be the intended config to be
      selected, but only a supporting config symbol to the actual config
      SND_SOC_WCD938X_SDW for the codec.

      Add SND_SOC_WCD938_SDW to the list instead of SND_SOC_WCD93{7,8}X.

      The issue was identified with ./scripts/checkkconfigsymbols.py.

      Fixes: 045442228868 ("ASoC: codecs: wcd938x: add audio routing and Kconfig")
      Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
      Link: https://lore.kernel.org/r/20211125095158.8394-3-lukas.bulwahn@gmail.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 1ad7c96b5a2986c150969d28198e7dda668cacd2
Author: Lukas Bulwahn <lukas.bulwahn@gmail.com>
Date:   Thu Nov 25 10:51:57 2021 +0100

    ASoC: uniphier: drop selecting non-existing SND_SOC_UNIPHIER_AIO_DMA

    [ Upstream commit 49f893253ab43566e34332a969324531fea463f6 ]

    Commit f37fe2f9987b ("ASoC: uniphier: add support for UniPhier AIO common
    driver") adds configs SND_SOC_UNIPHIER_{LD11,PXS2}, which select the
    non-existing config SND_SOC_UNIPHIER_AIO_DMA.

    Hence, ./scripts/checkkconfigsymbols.py warns:

      SND_SOC_UNIPHIER_AIO_DMA
      Referencing files: sound/soc/uniphier/Kconfig

    Probably, there is actually no further config intended to be selected
    here. So, just drop selecting the non-existing config.

    Fixes: f37fe2f9987b ("ASoC: uniphier: add support for UniPhier AIO common driver")
    Signed-off-by: Lukas Bulwahn <lukas.bulwahn@gmail.com>
    Link: https://lore.kernel.org/r/20211125095158.8394-2-lukas.bulwahn@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8a1f2ee9741a773ef6f5d600d39bb58f0fc691c
Author: Peiwei Hu <jlu.hpw@foxmail.com>
Date:   Fri Nov 19 17:12:18 2021 +0800

    powerpc/prom_init: Fix improper check of prom_getprop()

    [ Upstream commit 869fb7e5aecbc163003f93f36dcc26d0554319f6 ]

    prom_getprop() can return PROM_ERROR. Binary operator can not identify
    it.

    Fixes: 94d2dde738a5 ("[POWERPC] Efika: prune fixups and make them more carefull")
    Signed-off-by: Peiwei Hu <jlu.hpw@foxmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/tencent_BA28CC6897B7C95A92EB8C580B5D18589105@qq.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 98b0aba976790b9efe2879055e0018aff25a1a77
Author: Richard Fitzgerald <rf@opensource.cirrus.com>
Date:   Fri Nov 19 12:48:54 2021 +0000

    ASoC: cs42l42: Report initial jack state

    [ Upstream commit fdd535283779ec9f9c35fda352585c629121214f ]

    When a jack handler is registered in cs42l42_set_jack() the
    initial state should be reported if an attached headphone/headset
    has already been detected.

    The jack detect sequence takes around 1 second: typically long
    enough for the machine driver to probe and register the jack handler
    in time to receive the first report from the interrupt handler. So
    it is possible on some systems that the correct initial state was seen
    simply because of lucky timing. Modular builds were more likely to
    miss the reporting of the initial state.

    Signed-off-by: Richard Fitzgerald <rf@opensource.cirrus.com>
    Fixes: 4ca239f33737 ("ASoC: cs42l42: Always enable TS_PLUG and TS_UNPLUG interrupts")
```

```
    Link: https://lore.kernel.org/r/20211119124854.58939-1-rf@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef40bd07da2e7e17723afe9344801a60d3fef9c1
Author: Adam Ford <aford173@gmail.com>
Date:   Wed Nov 17 07:32:02 2021 -0600

    clk: imx8mn: Fix imx8mn_clko1_sels

    [ Upstream commit 570727e9acfac1c2330a01dd5e1272e9c3acec08 ]

    When attempting to use sys_pll1_80m as the parent for clko1, the
    system hangs.  This is due to the fact that the source select
    for sys_pll1_80m was incorrectly pointing to m7_alt_pll_clk, which
    doesn't yet exist.

    According to Rev 3 of the TRM, The imx8mn_clko1_sels also incorrectly
    references an osc_27m which does not exist, nor does an entry for
    source select bits 010b.  Fix both by inserting a dummy clock into
    the missing space in the table and renaming the incorrectly name clock
    with dummy.

    Fixes: 96d6392b54db ("clk: imx: Add support for i.MX8MN clock driver")
    Signed-off-by: Adam Ford <aford173@gmail.com>
    Reviewed-by: Fabio Estevam <festevam@gmail.com>
    Link: https://lore.kernel.org/r/20211117133202.775633-1-aford173@gmail.com
    Signed-off-by: Abel Vesa <abel.vesa@nxp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 463ea51b98e37f3e79f65cdfea67cf6e83165052
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
Date:   Wed Nov 17 11:50:59 2021 +0000

    clk: renesas: rzg2l: propagate return value of_genpd_add_provider_simple()

    [ Upstream commit 33748744f15a110a233b6ae0380f476006e770f0 ]

    of_genpd_add_provider_simple() might fail, this patch makes sure we check
    the return value of of_genpd_add_provider_simple() by propagating the
    return value to the caller of rzg2l_cpg_add_clk_domain().

    Fixes: ef3c613ccd68a ("clk: renesas: Add CPG core wrapper for RZ/G2L SoC")
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Link: https://lore.kernel.org/r/20211117115101.28281-3-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3efb62985c2f20652eb93c40ab92b364dd8524b5
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
Date:   Wed Nov 17 11:50:58 2021 +0000

    clk: renesas: rzg2l: Check return value of pm_genpd_init()

    [ Upstream commit 27527a3d3b162e4512798c058c0e8a216c721187 ]

    Make sure we check the return value of pm_genpd_init() which might fail.
    Also add a devres action to remove the power-domain in-case the probe
    callback fails further down in the code flow.

    Fixes: ef3c613ccd68a ("clk: renesas: Add CPG core wrapper for RZ/G2L SoC")
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Link: https://lore.kernel.org/r/20211117115101.28281-2-prabhakar.mahadev-
lad.rj@bp.renesas.com
```

```
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0462b0328fc20625ec40cb4433c6ba71446b553a
Author: Igor Pylypiv <ipylypiv@google.com>
Date:   Mon Nov 1 16:28:24 2021 -0700

    scsi: pm80xx: Update WARN_ON check in pm8001_mpi_build_cmd()

    [ Upstream commit 606c54ae975ad3af540b505b46b55a687501711f ]

    Starting from commit 05c6c029a44d ("scsi: pm80xx: Increase number of
    supported queues") driver initializes only max_q_num queues.  Do not use an
    invalid queue if the WARN_ON condition is true.

    Link: https://lore.kernel.org/r/20211101232825.2350233-4-ipylypiv@google.com
    Fixes: 7640e1eb8c5d ("scsi: pm80xx: Make mpi_build_cmd locking consistent")
    Reviewed-by: Vishakha Channapattan <vishakhavc@google.com>
    Acked-by: Jack Wang <jinpu.wang@ionos.com>
    Signed-off-by: Igor Pylypiv <ipylypiv@google.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e0b99a885d89eb1ff3d65f61d750577f337503b9
Author: Kamal Heib <kamalheib1@gmail.com>
Date:   Wed Nov 17 16:59:54 2021 +0200

    RDMA/hns: Validate the pkey index

    [ Upstream commit 2a67fcfa0db6b4075515bd23497750849b88850f ]

    Before query pkey, make sure that the queried index is valid.

    Fixes: 9a4435375cd1 ("IB/hns: Add driver files for hns RoCE driver")
    Link: https://lore.kernel.org/r/20211117145954.123893-1-kamalheib1@gmail.com
    Signed-off-by: Kamal Heib <kamalheib1@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f41e92cd54e773518e5c683865a169b4458bddbd
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Fri Nov 12 09:59:04 2021 +0100

    RDMA/bnxt_re: Scan the whole bitmap when checking if "disabling RCFW with pending cmd-
bit"

    [ Upstream commit a917dfb66c0a1fa1caacf3d71edcafcab48e6ff0 ]

    The 'cmdq->cmdq_bitmap' bitmap is 'rcfw->cmdq_depth' bits long.  The size
    stored in 'cmdq->bmap_size' is the size of the bitmap in bytes.

    Remove this erroneous 'bmap_size' and use 'rcfw->cmdq_depth' directly in
    'bnxt_qplib_disable_rcfw_channel()'. Otherwise some error messages may be
    missing.

    Other uses of 'cmdq_bitmap' already take into account 'rcfw->cmdq_depth'
    directly.

    Fixes: 1ac5a4047975 ("RDMA/bnxt_re: Add bnxt_re RoCE driver")
    Link:
https://lore.kernel.org/r/47ed717c3070a1d0f53e7b4c768a4fd11caf365d.1636707421.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Acked-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
```

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

commit 9910de4e54bc89c9906503a4d731e0066f1c12f0
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Nov 16 08:24:59 2021 +0100

    ALSA: hda: Fix potential deadlock at codec unbinding

    [ Upstream commit 7206998f578d5553989bc01ea2e544b622e79539 ]

    When a codec is unbound dynamically via sysfs while its stream is in
    use, we may face a potential deadlock at the proc remove or a UAF.
    This happens since the hda_pcm is managed by a linked list, as it
    handles the hda_pcm object release via kref.

    When a PCM is opened at the unbinding time, the release of hda_pcm
    gets delayed and it ends up with the close of the PCM stream releasing
    the associated hda_pcm object of its own.  The hda_pcm destructor
    contains the PCM device release that includes the removal of procfs
    entries.  And, this removal has the sync of the close of all in-use
    files -- which would never finish because it's called from the PCM
    file descriptor itself, i.e. it's trying to shoot its foot.

    For addressing the deadlock above, this patch changes the way to
    manage and release the hda_pcm object.  The kref of hda_pcm is
    dropped, and instead a simple refcount is introduced in hda_codec for
    keeping the track of the active PCM streams, and at each PCM open and
    close, this refcount is adjusted accordingly.  At unbinding, the
    driver calls snd_device_disconnect() for each PCM stream, then
    synchronizes with the refcount finish, and finally releases the object
    resources.

    Fixes: bbbc7e8502c9 ("ALSA: hda - Allocate hda_pcm objects dynamically")
    Link: https://lore.kernel.org/r/20211116072459.18930-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 39f05d6b16ca9a6c19c139918d0b83c816a893a8
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Nov 16 08:13:14 2021 +0100

    ALSA: hda: Add missing rwsem around snd_ctl_remove() calls

    [ Upstream commit 80bd64af75b4bb11c0329bc66c35da2ddfb66d88 ]

    snd_ctl_remove() has to be called with card->controls_rwsem held (when
    called after the card instantiation).  This patch add the missing
    rwsem calls around it.

    Fixes: d13bd412dce2 ("ALSA: hda - Manage kcontrol lists")
    Link: https://lore.kernel.org/r/20211116071314.15065-3-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 898499d56a59c363663067e68038ac9bed40acad
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Nov 16 08:13:13 2021 +0100

    ALSA: PCM: Add missing rwsem around snd_ctl_remove() calls

    [ Upstream commit 5471e9762e1af4b7df057a96bfd46cc250979b88 ]

    snd_ctl_remove() has to be called with card->controls_rwsem held (when
    called after the card instantiation).  This patch add the missing
    rwsem calls around it.
```

```
    Fixes: a8ff48cb7083 ("ALSA: pcm: Free chmap at PCM free callback, too")
    Link: https://lore.kernel.org/r/20211116071314.15065-2-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 849cb3f1e90a8f66e23f9f1e29d5da2171c81bd6
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Nov 16 08:13:12 2021 +0100

    ALSA: jack: Add missing rwsem around snd_ctl_remove() calls

    [ Upstream commit 06764dc931848c3a9bc01a63bbf76a605408bb54 ]

    snd_ctl_remove() has to be called with card->controls_rwsem held (when
    called after the card instantiation).  This patch add the missing
    rwsem calls around it.

    Fixes: 9058cbe1eed2 ("ALSA: jack: implement kctl creating for jack devices")
    Link: https://lore.kernel.org/r/20211116071314.15065-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f864cf32ed8d2cb1451ed00a942422faebdac0b3
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Wed Oct 27 10:18:16 2021 +0800

    ASoC: Intel: sof_sdw: fix jack detection on HP Spectre x360 convertible

    [ Upstream commit 0527b19fa4f390a6054612e1fa1dd4f8efc96739 ]

    Tests on device show the JD2 mode does not work at all, the 'Headphone
    Jack' and 'Headset Mic Jack' are shown as 'on' always.

    JD1 seems to be the better option, with at least a change between the
    two cases.

    Jack not plugged-in:
    [root@fedora ~]# amixer -Dhw:0 cget numid=12
    numid=12,iface=CARD,name='Headphone Jack'
      ; type=BOOLEAN,access=r-------,values=1
      : values=off
    [root@fedora ~]# amixer -Dhw:0 cget numid=13
    numid=13,iface=CARD,name='Headset Mic Jack'
      ; type=BOOLEAN,access=r-------,values=1
      : values=off

    Jack plugged-in:
    [root@fedora ~]# amixer -Dhw:0 cget numid=13
    numid=13,iface=CARD,name='Headset Mic Jack'
      ; type=BOOLEAN,access=r-------,values=1
      : values=on
    [root@fedora ~]# amixer -Dhw:0 cget numid=13
    numid=13,iface=CARD,name='Headset Mic Jack'
      ; type=BOOLEAN,access=r-------,values=1
      : values=on

    The 'Headset Mic Jack' is updated with a delay which seems normal with
    additional calibration needed.

    Fixes: d92e279dee56 ('ASoC: Intel: sof_sdw: add quirk for HP Spectre x360 convertible')
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20211027021824.24776-3-yung-chuan.liao@linux.intel.com
```

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

commit ab34f86e817a96daea4888d3a7e35d408c3b9699
Author: Jan Kara <jack@suse.cz>
Date:   Fri Nov 12 16:22:02 2021 +0100

    ext4: avoid trim error on fs with small groups

    [ Upstream commit 173b6e383d2a204c9921ffc1eca3b87aa2106c33 ]

    A user reported FITRIM ioctl failing for him on ext4 on some devices
    without apparent reason.  After some debugging we've found out that
    these devices (being LVM volumes) report rather large discard
    granularity of 42MB and the filesystem had 1k blocksize and thus group
    size of 8MB. Because ext4 FITRIM implementation puts discard
    granularity into minlen, ext4_trim_fs() declared the trim request as
    invalid. However just silently doing nothing seems to be a more
    appropriate reaction to such combination of parameters since user did
    not specify anything wrong.

    CC: Lukas Czerner <lczerner@redhat.com>
    Fixes: 5c2ed62fd447 ("ext4: Adjust minlen with discard_granularity in the FITRIM ioctl")
    Signed-off-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20211112152202.26614-1-jack@suse.cz
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 162dad32f4a6a8ff5e9d39504d4bffa8797a8f5b
Author: Taehee Yoo <ap420073@gmail.com>
Date:   Sun Jan 9 16:37:02 2022 +0000

    amt: fix wrong return type of amt_send_membership_update()

    [ Upstream commit dd3ca4c5184ea98e40acb8eb293d85b88ea04ee2 ]

    amt_send_membership_update() would return -1 but it's return type is bool.
    So, it should be used TRUE instead of -1.

    Fixes: cbc21dc1cfe9 ("amt: add data plane of amt interface")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Taehee Yoo <ap420073@gmail.com>
    Link: https://lore.kernel.org/r/20220109163702.6331-1-ap420073@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3622f49882a64dbbfa76dede080a990d6fc9543
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Fri Jan 7 01:57:16 2022 +0300

    net: mcs7830: handle usb read errors properly

    [ Upstream commit d668769eb9c52b150753f1653f7f5a0aeb8239d2 ]

    Syzbot reported uninit value in mcs7830_bind(). The problem was in
    missing validation check for bytes read via usbnet_read_cmd().

    usbnet_read_cmd() internally calls usb_control_msg(), that returns
    number of bytes read. Code should validate that requested number of bytes
    was actually read.

    So, this patch adds missing size validation check inside
    mcs7830_get_reg() to prevent uninit value bugs

    Reported-and-tested-by: syzbot+003c0a286b9af5412510@syzkaller.appspotmail.com
```

```
    Fixes: 2a36d7083438 ("USB: driver for mcs7830 (aka DeLOCK) USB ethernet adapter")
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20220106225716.7425-1-paskripkin@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e78fa8bab30c6aa2591471aaa523e3480f7936cc
Author: Edwin Peer <edwin.peer@broadcom.com>
Date:   Sun Jan 9 18:54:44 2022 -0500

    bnxt_en: use firmware provided max timeout for messages

    [ Upstream commit bce9a0b7900836df223ab638090df0cb8430d9e8 ]

    Some older devices cannot accommodate the 40 seconds timeout
    cap for long running commands (such as NVRAM commands) due to
    hardware limitations. Allow these devices to request more time for
    these long running commands, but print a warning, since the longer
    timeout may cause the hung task watchdog to trigger. In the case of a
    firmware update operation, this is preferable to failing outright.

    v2: Use bp->hwrm_cmd_max_timeout directly without the constants.

    Fixes: 881d8353b05e ("bnxt_en: Add an upper bound for all firmware command timeouts.")
    Signed-off-by: Edwin Peer <edwin.peer@broadcom.com>
    Signed-off-by: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a41cf0b9d6ab6caae9626116f7653131b1453a9f
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Mon Dec 27 12:17:57 2021 -0700

    iwlwifi: mvm: Use div_s64 instead of do_div in iwl_mvm_ftm_rtt_smoothing()

    [ Upstream commit 4ccdcc8ffd955490feec05380223db6a48961eb5 ]

    When building ARCH=arm allmodconfig:

    drivers/net/wireless/intel/iwlwifi/mvm/ftm-initiator.c: In function
‘iwl_mvm_ftm_rtt_smoothing’:
    ./include/asm-generic/div64.h:222:35: error: comparison of distinct pointer types lacks
a cast [-Werror]
      222 |         (void)(((typeof((n)) *)0) == ((uint64_t *)0));  \
          |                                   ^~
    drivers/net/wireless/intel/iwlwifi/mvm/ftm-initiator.c:1070:9: note: in expansion of
macro ‘do_div’
     1070 |         do_div(rtt_avg, 100);
          |         ^~~~~~

    do_div() has to be used with an unsigned 64-bit integer dividend but
    rtt_avg is a signed 64-bit integer.

    div_s64() expects a signed 64-bit integer dividend and signed 32-bit
    divisor, which fits this scenario, so use that function here to fix the
    warning.

    Fixes: 8b0f92549f2c ("iwlwifi: mvm: fix 32-bit build in FTM")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20211227191757.2354329-1-nathan@kernel.org
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 831e9fe7d746c890f5b13012d10368f6cc69bc1d
```

```
Author: Paul Blakey <paulb@nvidia.com>
Date:   Thu Jan 6 17:38:04 2022 +0200

    net: openvswitch: Fix ct_state nat flags for conns arriving from tc

    [ Upstream commit 6f022c2ddbcefaee79502ce5386dfe351d457070 ]

    Netfilter conntrack maintains NAT flags per connection indicating
    whether NAT was configured for the connection. Openvswitch maintains
    NAT flags on the per packet flow key ct_state field, indicating
    whether NAT was actually executed on the packet.

    When a packet misses from tc to ovs the conntrack NAT flags are set.
    However, NAT was not necessarily executed on the packet because the
    connection's state might still be in NEW state. As such, openvswitch
    wrongly assumes that NAT was executed and sets an incorrect flow key
    NAT flags.

    Fix this, by flagging to openvswitch which NAT was actually done in
    act_ct via tc_skb_ext and tc_skb_cb to the openvswitch module, so
    the packet flow key NAT flags will be correctly set.

    Fixes: b57dc7c13ea9 ("net/sched: Introduce action ct")
    Signed-off-by: Paul Blakey <paulb@nvidia.com>
    Acked-by: Jamal Hadi Salim <jhs@mojatatu.com>
    Link: https://lore.kernel.org/r/20220106153804.26451-1-paulb@nvidia.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edf221f69ae615b33244dae643b783512327c7d7
Author: Florian Westphal <fw@strlen.de>
Date:   Fri Jan 7 15:46:16 2022 +0100

    netfilter: egress: avoid a lockdep splat

    [ Upstream commit 6316136ec6e3dd1c302f7e7289a9ee46ecc610ae ]

    include/linux/netfilter_netdev.h:97 suspicious rcu_dereference_check() usage!
    2 locks held by sd-resolve/1100:
     0: ..(rcu_read_lock_bh){1:3}, at: ip_finish_output2
     1: ..(rcu_read_lock_bh){1:3}, at: __dev_queue_xmit
     __dev_queue_xmit+0 ..

    The helper has two callers, one uses rcu_read_lock, the other
    rcu_read_lock_bh().  Annotate the dereference to reflect this.

    Fixes: 42df6e1d221dd ("netfilter: Introduce egress hook")
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 329da166d261121e5fb78ab4cd11dfead23ccb96
Author: Dominik Brodowski <linux@dominikbrodowski.net>
Date:   Sun Jan 9 10:02:51 2022 +0100

    pcmcia: fix setting of kthread task states

    [ Upstream commit fbb3485f1f931102d8ba606f1c28123f5b48afa3 ]

    We need to set TASK_INTERRUPTIBLE before calling kthread_should_stop().
    Otherwise, kthread_stop() might see that the pccardd thread is still
    in TASK_RUNNING state and fail to wake it up.

    Additionally, we only need to set the state back to TASK_RUNNING if
    kthread_should_stop() breaks the loop.
```

```
    Cc: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Reported-by: Al Viro <viro@ZenIV.linux.org.uk>
    Reviewed-by: Matthew Wilcox (Oracle) <willy@infradead.org>
    Fixes: d3046ba809ce ("pcmcia: fix a boot time warning in pcmcia cs code")
    Signed-off-by: Dominik Brodowski <linux@dominikbrodowski.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be9a10ae2516f1cf4af743a520f01eb9282f1af3
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
Date:   Thu Jan 6 11:48:00 2022 +0000

    can: rcar_canfd: rcar_canfd_channel_probe(): make sure we free CAN network device

    [ Upstream commit 72b1e360572f9fa7d08ee554f1da29abce23f288 ]

    Make sure we free CAN network device in the error path. There are
    several jumps to fail label after allocating the CAN network device
    successfully. This patch places the free_candev() under fail label so
    that in failure path a jump to fail label frees the CAN network
    device.

    Fixes: 76e9353a80e9 ("can: rcar_canfd: Add support for RZ/G2L family")
    Link: https://lore.kernel.org/all/20220106114801.20563-1-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Reported-by: Pavel Machek <pavel@denx.de>
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Reviewed-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 72889036db13c2294d8efbc97e0aca2da7bb7cc0
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Fri Dec 24 10:13:24 2021 +0800

    can: xilinx_can: xcan_probe(): check for error irq

    [ Upstream commit c6564c13dae25cd7f8e1de5127b4da4500ee5844 ]

    For the possible failure of the platform_get_irq(), the returned irq
    could be error number and will finally cause the failure of the
    request_irq().

    Consider that platform_get_irq() can now in certain cases return
    -EPROBE_DEFER, and the consequences of letting request_irq()
    effectively convert that into -EINVAL, even at probe time rather than
    later on. So it might be better to check just now.

    Fixes: b1201e44f50b ("can: xilinx CAN controller support")
    Link: https://lore.kernel.org/all/20211224021324.1447494-1-jiasheng@iscas.ac.cn
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 749c023f31b9684baa266ea3f6d6c468143ff2f0
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Sat Jan 8 21:57:51 2022 +0100

    can: softing: softing_startstop(): fix set but not used variable warning

    [ Upstream commit 370d988cc529598ebaec6487d4f84c2115dc696b ]

    In the function softing_startstop() the variable error_reporting is
    assigned but not used. The code that uses this variable is commented
    out. Its stated that the functionality is not finally verified.
```

```
    To fix the warning:

    | drivers/net/can/softing/softing_fw.c:424:9: error: variable 'error_reporting' set but
 not used [-Werror,-Wunused-but-set-variable]

    remove the comment, activate the code, but add a "0 &&" to the if
    expression and rely on the optimizer rather than the preprocessor to
    remove the code.

    Link: https://lore.kernel.org/all/20220109103126.1872833-1-mkl@pengutronix.de
    Fixes: 03fd3cf5a179 ("can: add driver for Softing card")
    Cc: Kurt Van Dijck <dev.kurt@vandijck-laurijssen.be>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e58fa6bc6f905f7d4fdf3a479eeff76ebf76ee3
Author: Christophe Jaillet <christophe.jaillet@wanadoo.fr>
Date:   Sat Nov 6 17:42:04 2021 +0100

    tpm_tis: Fix an error handling path in 'tpm_tis_core_init()'

    [ Upstream commit e96d52822f5ac0a25de78f95cd23421bcbc93584 ]

    Commit 79ca6f74dae0 ("tpm: fix Atmel TPM crash caused by too frequent
    queries") has moved some code around without updating the error handling
    path.

    This is now pointless to 'goto out_err' when neither 'clk_enable()' nor
    'ioremap()' have been called yet.

    Make a direct return instead to avoid undoing things that have not been
    done.

    Fixes: 79ca6f74dae0 ("tpm: fix Atmel TPM crash caused by too frequent queries")
    Signed-off-by: Christophe Jaillet <christophe.jaillet@wanadoo.fr>
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 159610c96f14129f696256dd5056c76b2f3d1c0f
Author: Chen Jun <chenjun102@huawei.com>
Date:   Wed Oct 13 06:25:56 2021 +0000

    tpm: add request_locality before write TPM_INT_ENABLE

    [ Upstream commit 0ef333f5ba7f24f5d8478425c163d3097f1c7afd ]

    Locality is not appropriately requested before writing the int mask.
    Add the missing boilerplate.

    Fixes: e6aef069b6e9 ("tpm_tis: convert to using locality callbacks")
    Signed-off-by: Chen Jun <chenjun102@huawei.com>
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5bb90977a2ab3d0ead0293a0abadfe7ea9db76ca
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Tue Oct 19 17:05:25 2021 +0200

    can: mcp251xfd: add missing newline to printed strings

    [ Upstream commit 3bd9d8ce6f8c5c43ee2f1106021db0f98882cc75 ]
```

```
    This patch adds the missing newline to printed strings.

    Fixes: 55e5b97f003e ("can: mcp25xxfd: add driver for Microchip MCP25xxFD SPI CAN")
    Link: https://lore.kernel.org/all/20220105154300.1258636-4-mkl@pengutronix.de
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ab0fb180cfdd10701dd4b23f9ff5870291e26e0
Author: Sunil Goutham <sgoutham@marvell.com>
Date:   Fri Jan 7 12:25:05 2022 +0530

    octeontx2-af: Fix interrupt name strings

    [ Upstream commit 6dc9a23e29061e50c36523270de60039ccf536fa ]

    Fixed interrupt name string logic which currently results
    in wrong memory location being accessed while dumping
    /proc/interrupts.

    Fixes: 4826090719d4 ("octeontx2-af: Enable CPT HW interrupts")
    Signed-off-by: Sunil Goutham <sgoutham@marvell.com>
    Signed-off-by: Subbaraya Sundeep <sbhatta@marvell.com>
    Link: https://lore.kernel.org/r/1641538505-28367-1-git-send-email-sbhatta@marvell.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1261cbc5a109c00bbb4de8b094363fd63a3b6a4a
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Fri Jan 7 18:43:32 2022 +0200

    net: mscc: ocelot: fix incorrect balancing with down LAG ports

    [ Upstream commit a14e6b69f393d651913edcbe4ec0dec27b8b4b40 ]

    Assuming the test setup described here:
    https://patchwork.kernel.org/project/netdevbpf/cover/20210205130240.4072854-1-
vladimir.oltean@nxp.com/
    (swp1 and swp2 are in bond0, and bond0 is in a bridge with swp0)

    it can be seen that when swp1 goes down (on either board A or B), then
    traffic that should go through that port isn't forwarded anywhere.

    A dump of the PGID table shows the following:

    PGID_DST[0] = ports 0
    PGID_DST[1] = ports 1
    PGID_DST[2] = ports 2
    PGID_DST[3] = ports 3
    PGID_DST[4] = ports 4
    PGID_DST[5] = ports 5
    PGID_DST[6] = no ports
    PGID_AGGR[0] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[1] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[2] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[3] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[4] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[5] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[6] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[7] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[8] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[9] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[10] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[11] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[12] = ports 0, 1, 2, 3, 4, 5
    PGID_AGGR[13] = ports 0, 1, 2, 3, 4, 5
```

```
PGID_AGGR[14] = ports 0, 1, 2, 3, 4, 5
PGID_AGGR[15] = ports 0, 1, 2, 3, 4, 5
PGID_SRC[0] = ports 1, 2
PGID_SRC[1] = ports 0
PGID_SRC[2] = ports 0
PGID_SRC[3] = no ports
PGID_SRC[4] = no ports
PGID_SRC[5] = no ports
PGID_SRC[6] = ports 0, 1, 2, 3, 4, 5
```

Whereas a "good" PGID configuration for that setup should have looked like this:

```
PGID_DST[0] = ports 0
PGID_DST[1] = ports 1, 2
PGID_DST[2] = ports 1, 2
PGID_DST[3] = ports 3
PGID_DST[4] = ports 4
PGID_DST[5] = ports 5
PGID_DST[6] = no ports
PGID_AGGR[0] = ports 0, 2, 3, 4, 5
PGID_AGGR[1] = ports 0, 2, 3, 4, 5
PGID_AGGR[2] = ports 0, 2, 3, 4, 5
PGID_AGGR[3] = ports 0, 2, 3, 4, 5
PGID_AGGR[4] = ports 0, 2, 3, 4, 5
PGID_AGGR[5] = ports 0, 2, 3, 4, 5
PGID_AGGR[6] = ports 0, 2, 3, 4, 5
PGID_AGGR[7] = ports 0, 2, 3, 4, 5
PGID_AGGR[8] = ports 0, 2, 3, 4, 5
PGID_AGGR[9] = ports 0, 2, 3, 4, 5
PGID_AGGR[10] = ports 0, 2, 3, 4, 5
PGID_AGGR[11] = ports 0, 2, 3, 4, 5
PGID_AGGR[12] = ports 0, 2, 3, 4, 5
PGID_AGGR[13] = ports 0, 2, 3, 4, 5
PGID_AGGR[14] = ports 0, 2, 3, 4, 5
PGID_AGGR[15] = ports 0, 2, 3, 4, 5
PGID_SRC[0] = ports 1, 2
PGID_SRC[1] = ports 0
PGID_SRC[2] = ports 0
PGID_SRC[3] = no ports
PGID_SRC[4] = no ports
PGID_SRC[5] = no ports
PGID_SRC[6] = ports 0, 1, 2, 3, 4, 5
```

In other words, in the "bad" configuration, the attempt is to remove the inactive swp1 from the destination ports via PGID_DST. But when a MAC table entry is learned, it is learned towards PGID_DST 1, because that is the logical port id of the LAG itself (it is equal to the lowest numbered member port). So when swp1 becomes inactive, if we set PGID_DST[1] to contain just swp1 and not swp2, the packet will not have any chance to reach the destination via swp2.

The "correct" way to remove swp1 as a destination is via PGID_AGGR (remove swp1 from the aggregation port groups for all aggregation codes). This means that PGID_DST[1] and PGID_DST[2] must still contain both swp1 and swp2. This makes the MAC table still treat packets destined towards the single-port LAG as "multicast", and the inactive ports are removed via the aggregation code tables.

The change presented here is a design one: the ocelot_get_bond_mask() function used to take an "only_active_ports" argument. We don't need that. The only call site that specifies only_active_ports=true, ocelot_set_aggr_pgids(), must retrieve the entire bonding mask, because it must program that into PGID_DST. Additionally, it must also clear the inactive ports from the bond mask here, which it can't do if bond_mask

```
    just contains the active ports:

            ac = ocelot_read_rix(ocelot, ANA_PGID_PGID, i);
            ac &= ~bond_mask;  <---- here
            /* Don't do division by zero if there was no active
             * port. Just make all aggregation codes zero.
             */
            if (num_active_ports)
                    ac |= BIT(aggr_idx[i % num_active_ports]);
            ocelot_write_rix(ocelot, ac, ANA_PGID_PGID, i);

    So it becomes the responsibility of ocelot_set_aggr_pgids() to take
    ocelot_port->lag_tx_active into consideration when populating the
    aggr_idx array.

    Fixes: 23ca3b727ee6 ("net: mscc: ocelot: rebalance LAGs on link up/down events")
    Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Link: https://lore.kernel.org/r/20220107164332.402133-1-vladimir.oltean@nxp.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c230654deed9fabe2e1eee0907e01fd0ec0958ab
Author: Fabio Estevam <festevam@denx.de>
Date:   Fri Jan 7 13:33:07 2022 -0300

    regmap: Call regmap_debugfs_exit() prior to _init()

    [ Upstream commit 530792efa6cb86f5612ff093333fec735793b582 ]

    Since commit cffa4b2122f5 ("regmap: debugfs: Fix a memory leak when
    calling regmap_attach_dev"), the following debugfs error is seen
    on i.MX boards:

    debugfs: Directory 'dummy-iomuxc-gpr@20e0000' with parent 'regmap' already present!

    In the attempt to fix the memory leak, the above commit added a NULL check
    for map->debugfs_name. For the first debufs entry, map->debugfs_name is NULL
    and then the new name is allocated via kasprintf().

    For the second debugfs entry, map->debugfs_name() is no longer NULL, so
    it will keep using the old entry name and the duplicate name error is seen.

    Quoting Mark Brown:

    "That means that if the device gets freed we'll end up with the old debugfs
    file hanging around pointing at nothing.
    ...
    To be more explicit this means we need a call to regmap_debugfs_exit()
    which will clean up all the existing debugfs stuff before we loose
    references to it."

    Call regmap_debugfs_exit() prior to regmap_debugfs_init() to fix
    the problem.

    Tested on i.MX6Q and i.MX6SX boards.

    Fixes: cffa4b2122f5 ("regmap: debugfs: Fix a memory leak when calling
regmap_attach_dev")
    Suggested-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Fabio Estevam <festevam@denx.de>
    Link: https://lore.kernel.org/r/20220107163307.335404-1-festevam@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 25dff18a1e8db244289cbe2dc49345d8440863a6
```

```
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jan 7 10:12:10 2022 +0300

    netrom: fix api breakage in nr_setsockopt()

    [ Upstream commit dc35616e6c2907b0c0c391a205802d8880f7fd85 ]

    This needs to copy an unsigned int from user space instead of a long to
    avoid breaking user space with an API change.

    I have updated all the integer overflow checks from ULONG to UINT as
    well.  This is a slight API change but I do not expect it to affect
    anything in real life.

    Fixes: 3087a6f36ee0 ("netrom: fix copying in user data in nr_setsockopt")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9cdfbf4bc7bbf8d237c987d7c1da05fc0d7e869c
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jan 7 10:13:12 2022 +0300

    ax25: uninitialized variable in ax25_setsockopt()

    [ Upstream commit 9371937092d5fd502032c1bb4475b36b39b1f1b3 ]

    The "opt" variable is unsigned long but we only copy 4 bytes from
    the user so the lower 4 bytes are uninitialized.

    I have changed the integer overflow checks from ULONG to UINT as well.
    This is a slight API change but I don't expect it to break anything.

    Fixes: a7b75c5a8c41 ("net: pass a sockptr_t into ->setsockopt")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b03758e5c424784450e38db50d8d221ffa7afdb
Author: Rakesh Babu Saladi <rsaladi2@marvell.com>
Date:   Fri Jan 7 12:00:30 2022 +0530

    octeontx2-nicvf: Free VF PTP resources.

    [ Upstream commit eabd0f88b0d2d433c5dfe88218d4ce1c11ef04b8 ]

    When a VF is removed respective PTP resources are not
    being freed currently. This patch fixes it.

    Fixes: 43510ef4ddad ("octeontx2-nicvf: Add PTP hardware clock support to NIX VF")
    Signed-off-by: Rakesh Babu Saladi <rsaladi2@marvell.com>
    Signed-off-by: Subbaraya Sundeep <sbhatta@marvell.com>
    Signed-off-by: Sunil Goutham <sgoutham@marvell.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 091917a5ba55b0a1281df9e94301068f6cb3caa3
Author: Subbaraya Sundeep <sbhatta@marvell.com>
Date:   Fri Jan 7 12:00:29 2022 +0530

    octeontx2-af: Increment ptp refcount before use

    [ Upstream commit 93440f4888cf049dbd22b41aaf94d2e2153b3eb8 ]

    Before using the ptp pci device by AF driver increment
```

```
     the reference count of it.

     Fixes: a8b90c9d26d6 ("octeontx2-af: Add PTP device id for CN10K and 95O silcons")
     Signed-off-by: Subbaraya Sundeep <sbhatta@marvell.com>
     Signed-off-by: Sunil Goutham <sgoutham@marvell.com>
     Signed-off-by: David S. Miller <davem@davemloft.net>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit ab9e776091b722780bc9fc76c5d93cb01572eae6
 Author: Miaoqian Lin <linmq006@gmail.com>
 Date:   Fri Jan 7 07:54:24 2022 +0000

     spi: spi-meson-spifc: Add missing pm_runtime_disable() in meson_spifc_probe

     [ Upstream commit 69c1b87516e327a60b39f96b778fe683259408bf ]

     If the probe fails, we should use pm_runtime_disable() to balance
     pm_runtime_enable().
     Add missing pm_runtime_disable() for meson_spifc_probe.

     Fixes: c3e4bc5434d2 ("spi: meson: Add support for Amlogic Meson SPIFC")
     Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
     Link: https://lore.kernel.org/r/20220107075424.7774-1-linmq006@gmail.com
     Signed-off-by: Mark Brown <broonie@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 26e60c0ab6313a5d31619242380d13621d141431
 Author: Mat Martineau <mathew.j.martineau@linux.intel.com>
 Date:   Thu Jan 6 14:06:38 2022 -0800

     mptcp: Check reclaim amount before reducing allocation

     [ Upstream commit 269bda9e7da48eafb599d01c96199caa2f7547e5 ]

     syzbot found a page counter underflow that was triggered by MPTCP's
     reclaim code:

     page_counter underflow: -4294964789 nr_pages=4294967295
     WARNING: CPU: 2 PID: 3785 at mm/page_counter.c:56 page_counter_cancel+0xcf/0xe0
 mm/page_counter.c:56
     Modules linked in:
     CPU: 2 PID: 3785 Comm: kworker/2:6 Not tainted 5.16.0-rc1-syzkaller #0
     Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.14.0-2 04/01/2014
     Workqueue: events mptcp_worker

     RIP: 0010:page_counter_cancel+0xcf/0xe0 mm/page_counter.c:56
     Code: c7 04 24 00 00 00 00 45 31 f6 eb 97 e8 2a 2b b5 ff 4c 89 ea 48 89 ee 48 c7 c7 00
 9e b8 89 c6 05 a0 c1 ba 0b 01 e8 95 e4 4b 07 <0f> 0b eb a8 4c 89 e7 e8 25 5a fb ff eb c7 0f
 1f 00 41 56 41 55 49
     RSP: 0018:ffffc90002d4f918 EFLAGS: 00010082

     RAX: 0000000000000000 RBX: ffff88806a494120 RCX: 0000000000000000
     RDX: ffff8880688c41c0 RSI: ffffffff815e8f28 RDI: fffff520005a9f15
     RBP: ffffffff000009cb R08: 0000000000000000 R09: 0000000000000000
     R10: ffffffff815e2cfe R11: 0000000000000000 R12: ffff88806a494120
     R13: 00000000ffffffff R14: 0000000000000000 R15: 0000000000000001
     FS:  0000000000000000(0000) GS:ffff88802cc00000(0000) knlGS:0000000000000000
     CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
     CR2: 0000001b2de21000 CR3: 000000005ad59000 CR4: 0000000000150ee0
     DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
     DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
     Call Trace:
      <TASK>
      page_counter_uncharge+0x2e/0x60 mm/page_counter.c:160
      drain_stock+0xc1/0x180 mm/memcontrol.c:2219
```

```
       refill_stock+0x139/0x2f0 mm/memcontrol.c:2271
       __sk_mem_reduce_allocated+0x24d/0x550 net/core/sock.c:2945
       __mptcp_rmem_reclaim net/mptcp/protocol.c:167 [inline]
       __mptcp_mem_reclaim_partial+0x124/0x410 net/mptcp/protocol.c:975
       mptcp_mem_reclaim_partial net/mptcp/protocol.c:982 [inline]
       mptcp_alloc_tx_skb net/mptcp/protocol.c:1212 [inline]
       mptcp_sendmsg_frag+0x18c6/0x2190 net/mptcp/protocol.c:1279
       __mptcp_push_pending+0x232/0x720 net/mptcp/protocol.c:1545
       mptcp_release_cb+0xfe/0x200 net/mptcp/protocol.c:2975
       release_sock+0xb4/0x1b0 net/core/sock.c:3306
       mptcp_worker+0x51e/0xc10 net/mptcp/protocol.c:2443
       process_one_work+0x9b2/0x1690 kernel/workqueue.c:2298
       worker_thread+0x658/0x11f0 kernel/workqueue.c:2445
       kthread+0x405/0x4f0 kernel/kthread.c:327
       ret_from_fork+0x1f/0x30 arch/x86/entry/entry_64.S:295
       </TASK>

    __mptcp_mem_reclaim_partial() could call __mptcp_rmem_reclaim() with a
    negative value, which passed that negative value to
    __sk_mem_reduce_allocated() and triggered the splat above.

    Check for a reclaim amount that is positive and large enough for
    __mptcp_rmem_reclaim() to actually adjust rmem_fwd_alloc (much like
    the sk_mem_reclaim_partial() code the function is based on).

    v2: Use '>' instead of '>=', since SK_MEM_QUANTUM - 1 would get
    right-shifted into nothing by __mptcp_rmem_reclaim.

    Fixes: 6511882cdd82 ("mptcp: allocate fwd memory separately on the rx and tx path")
    Closes: https://github.com/multipath-tcp/mptcp_net-next/issues/252
    Reported-and-tested-by: syzbot+bc9e2d2dbcb347dd215a@syzkaller.appspotmail.com
    Cc: Andrew Morton <akpm@linux-foundation.org>
    Cc: Michal Hocko <mhocko@suse.com>
    Acked-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef50378f2e687bb779aa1382e65b36b0b384eebd
Author: Geliang Tang <geliang.tang@suse.com>
Date:   Thu Jan 6 14:06:37 2022 -0800

    mptcp: fix a DSS option writing error

    [ Upstream commit 110b6d1fe98fd7af9893992459b651594d789293 ]

    'ptr += 1;' was omitted in the original code.

    If the DSS is the last option -- which is what we have most of the
    time -- that's not an issue. But it is if we need to send something else
    after like a RM_ADDR or an MP_PRIO.

    Fixes: 1bff1e43a30e ("mptcp: optimize out option generation")
    Reviewed-by: Matthieu Baerts <matthieu.baerts@tessares.net>
    Signed-off-by: Geliang Tang <geliang.tang@suse.com>
    Signed-off-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68debee5bf2fa0b104b81d4910cc9c844fa65ca1
Author: Matthieu Baerts <matthieu.baerts@tessares.net>
Date:   Thu Jan 6 14:06:36 2022 -0800

    mptcp: fix opt size when sending DSS + MP_FAIL
```

```
    [ Upstream commit 04fac2cae9422a3401c172571afbcfdd58fa5c7e ]

    When these two options had to be sent -- which is not common -- the DSS
    size was not being taken into account in the remaining size.

    Additionally in this situation, the reported size was only the one of
    the MP_FAIL which can cause issue if at the end, we need to write more
    in the TCP options than previously said.

    Here we use a dedicated variable for MP_FAIL size to keep the
    WARN_ON_ONCE() just after.

    Fixes: c25aeb4e0953 ("mptcp: MP_FAIL suboption sending")
    Acked-and-tested-by: Geliang Tang <geliang.tang@suse.com>
    Signed-off-by: Matthieu Baerts <matthieu.baerts@tessares.net>
    Signed-off-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3a9a1e79fd6b0c56f725e193669b928c7d6089e
Author: Paolo Abeni <pabeni@redhat.com>
Date:   Thu Jan 6 16:20:20 2022 -0800

    mptcp: fix per socket endpoint accounting

    [ Upstream commit f7d6a237d7422809d458d754016de2844017cb4d ]

    Since full-mesh endpoint support, the reception of a single ADD_ADDR
    option can cause multiple subflows creation. When such option is
    accepted we increment 'add_addr_accepted' by one. When we received
    a paired RM_ADDR option, we deleted all the relevant subflows,
    decrementing 'add_addr_accepted' by one for each of them.

    We have a similar issue for 'local_addr_used'

    Fix them moving the pm endpoint accounting outside the subflow
    traversal.

    Fixes: 1a0d6136c5f0 ("mptcp: local addresses fullmesh")
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Mat Martineau <mathew.j.martineau@linux.intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f8e43c8314cca5b8b73dcc111cb977f8918acfd
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jan 7 10:17:27 2022 +0300

    Bluetooth: hci_sock: fix endian bug in hci_sock_setsockopt()

    [ Upstream commit b9f9dbad0bd1c302d357fdd327c398f51f5fc2b1 ]

    This copies a u16 into the high bits of an int, which works on a big
    endian system but not on a little endian system.

    Fixes: 09572fca7223 ("Bluetooth: hci_sock: Add support for BT_{SND,RCV}BUF")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8bf7c5afabe3354a6b5ca7c2108aa6a289b11aeb
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Jan 7 10:16:44 2022 +0300

    Bluetooth: L2CAP: uninitialized variables in l2cap_sock_setsockopt()
```

```
    [ Upstream commit 2b70d4f9b20635ac328836e50d183632e1930f94 ]

    The "opt" variable is a u32, but on some paths only the top bytes
    were initialized and the others contained random stack data.

    Fixes: a7b75c5a8c41 ("net: pass a sockptr_t into ->setsockopt")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f23826a62d392bc133fa46a1fb6655e5ac8270cb
Author: Zizhuang Deng <sunsetdzz@gmail.com>
Date:   Thu Dec 30 15:03:31 2021 +0800

    lib/mpi: Add the return value check of kcalloc()

    [ Upstream commit dd827abe296fe4249b2f8c9b95f72f814ea8348c ]

    Add the return value check of kcalloc() to avoid potential
    NULL ptr dereference.

    Fixes: a8ea8bdd9df9 ("lib/mpi: Extend the MPI library")
    Signed-off-by: Zizhuang Deng <sunsetdzz@gmail.com>
    Reviewed-by: Tianjia Zhang <tianjia.zhang@linux.alibaba.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4646175c19fd019b773444a11ff62748eb83745b
Author: Moshe Shemesh <moshe@nvidia.com>
Date:   Sun Dec 5 12:07:49 2021 +0200

    net/mlx5: Set command entry semaphore up once got index free

    [ Upstream commit 8e715cd613a1e872b9d918e912d90b399785761a ]

    Avoid a race where command work handler may fail to allocate command
    entry index, by holding the command semaphore down till command entry
    index is being freed.

    Fixes: 410bd754cd73 ("net/mlx5: Add retry mechanism to the command entry index
allocation")
    Signed-off-by: Moshe Shemesh <moshe@nvidia.com>
    Reviewed-by: Eran Ben Elisha <eranbe@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edcd9d1cb405f4740d803b2e78ad11025dfdc6c1
Author: Maor Dickman <maord@nvidia.com>
Date:   Mon Jan 3 15:04:18 2022 +0200

    net/mlx5e: Sync VXLAN udp ports during uplink representor profile change

    [ Upstream commit 07f6dc4024ea1d2314b9c8b81fd4e492864fcca1 ]

    Currently during NIC profile disablement all VXLAN udp ports offloaded to the
    HW are flushed and during its enablement the driver send notification to
    the stack to inform the core that the entire UDP tunnel port state has been
    lost, uplink representor doesn't have the same behavior which can cause
    VXLAN udp ports offload to be in bad state while moving between modes while
    VXLAN interface exist.

    Fixed by aligning the uplink representor profile behavior to the NIC behavior.

    Fixes: 84db66124714 ("net/mlx5e: Move set vxlan nic info to profile init")
```

```
    Signed-off-by: Maor Dickman <maord@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f1952b0abeb6de85ef87dd550a5cdc48a5ca595
Author: Shay Drory <shayd@nvidia.com>
Date:   Thu Dec 30 08:54:08 2021 +0200

    net/mlx5: Fix access to sf_dev_table on allocation failure

    [ Upstream commit a1c7c49c2091926962f8c1c866d386febffec5d8 ]

    Even when SF devices are supported, the SF device table allocation
    can still fail.
    In such case mlx5_sf_dev_supported still reports true, but SF device
    table is invalid. This can result in NULL table access.

    Hence, fix it by adding NULL table check.

    Fixes: 1958fc2f0712 ("net/mlx5: SF, Add auxiliary device driver")
    Signed-off-by: Shay Drory <shayd@nvidia.com>
    Reviewed-by: Parav Pandit <parav@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 040329285f110b1d4d78d5dcc111e4de2d5ee18f
Author: Paul Blakey <paulb@nvidia.com>
Date:   Wed Jun 16 10:55:56 2021 +0300

    net/mlx5e: Fix matching on modified inner ip_ecn bits

    [ Upstream commit b6dfff21a170af5c695ebaa153b7f5e297ddca03 ]

    Tunnel device follows RFC 6040, and during decapsulation inner
    ip_ecn might change depending on inner and outer ip_ecn as follows:

     +---------+----------------------------------------+
     |Arriving |         Arriving Outer Header          |
     |   Inner +---------+---------+---------+----------+
     |  Header | Not-ECT | ECT(0)  | ECT(1)  |   CE     |
     +---------+---------+---------+---------+----------+
     | Not-ECT | Not-ECT | Not-ECT | Not-ECT | <drop>   |
     |  ECT(0) |  ECT(0) | ECT(0)  | ECT(1)  |   CE*    |
     |  ECT(1) |  ECT(1) | ECT(1)  | ECT(1)* |   CE*    |
     |    CE   |   CE    |  CE     | CE      |   CE     |
     +---------+---------+---------+---------+----------+

    Cells marked above are changed from original inner packet ip_ecn value.

    Tc then matches on the modified inner ip_ecn, but hw offload which
    matches the inner ip_ecn value before decap, will fail.

    Fix that by mapping all the cases of outer and inner ip_ecn matching,
    and only supporting cases where we know inner wouldn't be changed by
    decap, or in the outer ip_ecn=CE case, inner ip_ecn didn't matter.

    Fixes: bcef735c59f2 ("net/mlx5e: Offload TC matching on tos/ttl for ip tunnels")
    Signed-off-by: Paul Blakey <paulb@nvidia.com>
    Reviewed-by: Oz Shlomo <ozsh@nvidia.com>
    Reviewed-by: Eli Cohen <elic@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 793641395e7c1bdd6b1c9d96d6e38dc9a7f9eff0
Author: Aya Levin <ayal@nvidia.com>
Date:   Sun Oct 24 16:52:23 2021 +0300

    Revert "net/mlx5e: Block offload of outer header csum for GRE tunnel"

    [ Upstream commit 01c3fd113ef50490ffd43f78f347ef6bb008510b ]

    This reverts commit 54e1217b90486c94b26f24dcee1ee5ef5372f832.

    Although the NIC doesn't support offload of outer header CSUM, using
    gso_partial_features allows offloading the tunnel's segmentation. The
    driver relies on the stack CSUM calculation of the outer header. For
    this, NETIF_F_GSO_GRE_CSUM must be a member of the device's features.

    Fixes: 54e1217b9048 ("net/mlx5e: Block offload of outer header csum for GRE tunnel")
    Signed-off-by: Aya Levin <ayal@nvidia.com>
    Reviewed-by: Gal Pressman <gal@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6a85f3b161f0b9e33f17b8b2d9846e57a7d75761
Author: Aya Levin <ayal@nvidia.com>
Date:   Sun Oct 24 11:47:41 2021 +0300

    Revert "net/mlx5e: Block offload of outer header csum for UDP tunnels"

    [ Upstream commit 64050cdad0983ad8060e33c3f4b5aee2366bcebd ]

    This reverts commit 6d6727dddc7f93fcc155cb8d0c49c29ae0e71122.

    Although the NIC doesn't support offload of outer header CSUM, using
    gso_partial_features allows offloading the tunnel's segmentation. The
    driver relies on the stack CSUM calculation of the outer header. For
    this, NETIF_F_GSO_UDP_TUNNEL_CSUM must be a member of the device's
    features.

    Fixes: 6d6727dddc7f ("net/mlx5e: Block offload of outer header csum for UDP tunnels")
    Signed-off-by: Aya Levin <ayal@nvidia.com>
    Reviewed-by: Gal Pressman <gal@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce848c093ecc01a9d804565fac8cf07063cab3ba
Author: Maor Dickman <maord@nvidia.com>
Date:   Thu Dec 30 11:20:10 2021 +0200

    net/mlx5e: Don't block routes with nexthop objects in SW

    [ Upstream commit 9e72a55a3c9d54b38a704bb7292d984574a81d9d ]

    Routes with nexthop objects is currently not supported by multipath offload
    and any attempts to use it is blocked, however this also block adding SW
    routes with nexthop.

    Resolve this by returning NOTIFY_DONE instead of an error which will allow such
    a route to be created in SW but not offloaded.

    This fix also solve an issue which block adding such routes on different devices
    due to missing check if the route FIB device is one of multipath devices.

    Fixes: 6a87afc072c3 ("mlx5: Fail attempts to use routes with nexthop objects")
    Signed-off-by: Maor Dickman <maord@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
```

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

commit 1230da0a583408ad4507ea6e61ec0105e7ad31a2
Author: Maor Dickman <maord@nvidia.com>
Date:   Wed Dec 29 16:10:41 2021 +0200

    net/mlx5e: Fix wrong usage of fib_info_nh when routes with nexthop objects are used

    [ Upstream commit 885751eb1b01d276e38f57d78c583e4ce006c5ed ]

    Creating routes with nexthop objects while in switchdev mode leads to access to
    un-allocated memory and trigger bellow call trace due to hitting WARN_ON.
    This is caused due to illegal usage of fib_info_nh in TC tunnel FIB event handling to
    resolve the FIB device while fib_info built in with nexthop.

    Fixed by ignoring attempts to use nexthop objects with routes until support can be
    properly added.

    WARNING: CPU: 1 PID: 1724 at include/net/nexthop.h:468
mlx5e_tc_tun_fib_event+0x448/0x570 [mlx5_core]
    CPU: 1 PID: 1724 Comm: ip Not tainted 5.15.0_for_upstream_min_debug_2021_11_09_02_04 #1
    RIP: 0010:mlx5e_tc_tun_fib_event+0x448/0x570 [mlx5_core]
    RSP: 0018:ffff8881349f7910 EFLAGS: 00010202
    RAX: ffff8881492f1980 RBX: ffff8881349f79e8 RCX: 0000000000000000
    RDX: ffff8881349f79e8 RSI: 0000000000000000 RDI: 0000000000000000
    RBP: ffff8881349f7950 R08: 00000000000000fe R09: 0000000000000001
    R10: 0000000000000000 R11: 0000000000000000 R12: ffff88811e9d0000
    R13: ffff88810eb62000 R14: ffff888106710268 R15: 0000000000000018
    FS:  00007f1d5ca6e800(0000) GS:ffff88852c880000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00007ffedba44ff8 CR3: 0000000129808004 CR4: 0000000000370ea0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     <TASK>
     atomic_notifier_call_chain+0x42/0x60
     call_fib_notifiers+0x21/0x40
     fib_table_insert+0x479/0x6d0
     ? try_charge_memcg+0x480/0x6d0
     inet_rtm_newroute+0x65/0xb0
     rtnetlink_rcv_msg+0x2af/0x360
     ? page_add_file_rmap+0x13/0x130
     ? do_set_pte+0xcd/0x120
     ? rtnl_calcit.isra.0+0x120/0x120
     netlink_rcv_skb+0x4e/0xf0
     netlink_unicast+0x1ee/0x2b0
     netlink_sendmsg+0x22e/0x460
     sock_sendmsg+0x33/0x40
     ____sys_sendmsg+0x1d1/0x1f0
     ___sys_sendmsg+0xab/0xf0
     ? __mod_memcg_lruvec_state+0x40/0x60
     ? __mod_lruvec_page_state+0x95/0xd0
     ? page_add_new_anon_rmap+0x4e/0xf0
     ? __handle_mm_fault+0xec6/0x1470
     __sys_sendmsg+0x51/0x90
     ? internal_get_user_pages_fast+0x480/0xa10
     do_syscall_64+0x3d/0x90
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    Fixes: 8914add2c9e5 ("net/mlx5e: Handle FIB events to update tunnel endpoint device")
    Signed-off-by: Maor Dickman <maord@nvidia.com>
    Reviewed-by: Vlad Buslov <vladbu@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 6ba7ef1391bba8cc050c8a1e829d991061238d49
Author: Dima Chumak <dchumak@nvidia.com>
Date:   Wed Nov 24 09:37:26 2021 +0200

    net/mlx5e: Fix nullptr on deleting mirroring rule

    [ Upstream commit de31854ece175e12ff3c35d07f340988823aed34 ]

    Deleting a Tc rule with multiple outputs, one of which is internal port,
    like this one:

      tc filter del dev enp8s0f0_0 ingress protocol ip pref 5 flower \
          dst_mac 0c:42:a1:d1:d0:88 \
          src_mac e4:ea:09:08:00:02 \
          action tunnel_key  set \
              src_ip 0.0.0.0 \
              dst_ip 7.7.7.8 \
              id 8 \
              dst_port 4789 \
          action mirred egress mirror dev vxlan_sys_4789 pipe \
          action mirred egress redirect dev enp8s0f0_1

    Triggers a call trace:

      BUG: kernel NULL pointer dereference, address: 0000000000000230
      RIP: 0010:del_sw_hw_rule+0x2b/0x1f0 [mlx5_core]
      Call Trace:
       tree_remove_node+0x16/0x30 [mlx5_core]
       mlx5_del_flow_rules+0x51/0x160 [mlx5_core]
       __mlx5_eswitch_del_rule+0x4b/0x170 [mlx5_core]
       mlx5e_tc_del_fdb_flow+0x295/0x550 [mlx5_core]
       mlx5e_flow_put+0x1f/0x70 [mlx5_core]
       mlx5e_delete_flower+0x286/0x390 [mlx5_core]
       tc_setup_cb_destroy+0xac/0x170
       fl_hw_destroy_filter+0x94/0xc0 [cls_flower]
       __fl_delete+0x15e/0x170 [cls_flower]
       fl_delete+0x36/0x80 [cls_flower]
       tc_del_tfilter+0x3a6/0x6e0
       rtnetlink_rcv_msg+0xe5/0x360
       ? rtnl_calcit.isra.0+0x110/0x110
       netlink_rcv_skb+0x46/0x110
       netlink_unicast+0x16b/0x200
       netlink_sendmsg+0x202/0x3d0
       sock_sendmsg+0x33/0x40
       ____sys_sendmsg+0x1c3/0x200
       ? copy_msghdr_from_user+0xd6/0x150
       ___sys_sendmsg+0x88/0xd0
       ? ___sys_recvmsg+0x88/0xc0
       ? do_futex+0x10c/0x460
       __sys_sendmsg+0x59/0xa0
       do_syscall_64+0x48/0x140
       entry_SYSCALL_64_after_hwframe+0x44/0xa9

    Fix by disabling offloading for flows matching
    esw_is_chain_src_port_rewrite() which have more than one output.

    Fixes: 10742efc20a4 ("net/mlx5e: VF tunnel TX traffic offloading")
    Signed-off-by: Dima Chumak <dchumak@nvidia.com>
    Reviewed-by: Roi Dayan <roid@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9e34091c4366b646331939583745d9dbf44b4338
Author: Aya Levin <ayal@nvidia.com>
```

```
Date:   Thu Dec 23 14:38:28 2021 +0200

    net/mlx5e: Fix page DMA map/unmap attributes

    [ Upstream commit 0b7cfa4082fbf550595bc0e40f05614bd83bf0cd ]

    Driver initiates DMA sync, hence it may skip CPU sync. Add
    DMA_ATTR_SKIP_CPU_SYNC as input attribute both to dma_map_page and
    dma_unmap_page to avoid redundant sync with the CPU.
    When forcing the device to work with SWIOTLB, the extra sync might cause
    data corruption. The driver unmaps the whole page while the hardware
    used just a part of the bounce buffer. So syncing overrides the entire
    page with bounce buffer that only partially contains real data.

    Fixes: bc77b240b3c5 ("net/mlx5e: Add fragmented memory support for RX multi packet WQE")
    Fixes: db05815b36cb ("net/mlx5e: Add XSK zero-copy support")
    Signed-off-by: Aya Levin <ayal@nvidia.com>
    Reviewed-by: Gal Pressman <gal@nvidia.com>
    Signed-off-by: Saeed Mahameed <saeedm@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a441229e91d3f07813967e9ef020badb979c03ca
Author: Huang Rui <ray.huang@amd.com>
Date:   Thu Jan 6 15:43:06 2022 +0800

    x86, sched: Fix undefined reference to init_freq_invariance_cppc() build error

    [ Upstream commit 6c4ab1b86dac3954d15c00c1a6396d60a1023fab ]

    The init_freq_invariance_cppc function is implemented in smpboot and depends on
    CONFIG_SMP.

      MODPOST vmlinux.symvers
      MODINFO modules.builtin.modinfo
      GEN     modules.builtin
      LD      .tmp_vmlinux.kallsyms1
    ld: drivers/acpi/cppc_acpi.o: in function `acpi_cppc_processor_probe':
    /home/ray/brahma3/linux/drivers/acpi/cppc_acpi.c:819: undefined reference to
`init_freq_invariance_cppc'
    make: *** [Makefile:1161: vmlinux] Error 1

    See https://lore.kernel.org/lkml/484af487-7511-647e-5c5b-33d4429acdec@infradead.org/.

    Fixes: 41ea667227ba ("x86, sched: Calculate frequency invariance for AMD systems")
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Randy Dunlap <rdunlap@infradead.org>
    Reported-by: Stephen Rothwell <sfr@canb.auug.org.au>
    Signed-off-by: Huang Rui <ray.huang@amd.com>
    [ rjw: Subject edits ]
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ed6448ef084bd2c911b5c519f93c140a5c0c901
Author: Vinod Koul <vkoul@kernel.org>
Date:   Mon Jan 3 12:41:18 2022 +0530

    spi: qcom: geni: handle timeout for gpi mode

    [ Upstream commit f8039ea55d4ccac2238a247a574f0acb3bc1dc4b ]

    We missed adding handle_err for gpi mode, so add a new function
    spi_geni_handle_err() which would call handle_fifo_timeout() or newly
    added handle_gpi_timeout() based on mode

    Fixes: b59c122484ec ("spi: spi-geni-qcom: Add support for GPI dma")
```

```
    Reported-by: Douglas Anderson <dianders@chromium.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Link: https://lore.kernel.org/r/20220103071118.27220-2-vkoul@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 95b2871c9d1f75b8ea8a3da03aeef6c0a35036ad
Author: Vinod Koul <vkoul@kernel.org>
Date:   Mon Jan 3 12:41:17 2022 +0530

    spi: qcom: geni: set the error code for gpi transfer

    [ Upstream commit 74b86d6af81be73bb74995ebeba74417e84b6b6f ]

    Before we invoke spi_finalize_current_transfer() in
    spi_gsi_callback_result() we should set the spi->cur_msg->status as
    appropriate (0 for success, error otherwise).

    The helps to return error on transfer and not wait till it timesout on
    error

    Fixes: b59c122484ec ("spi: spi-geni-qcom: Add support for GPI dma")
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Link: https://lore.kernel.org/r/20220103071118.27220-1-vkoul@kernel.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8a81859aab32d2cfd5c0d3cfc3ea18c4b95355e
Author: Valentin Caron <valentin.caron@foss.st.com>
Date:   Tue Jan 4 19:24:42 2022 +0100

    serial: stm32: move tx dma terminate DMA to shutdown

    [ Upstream commit 56a23f9319e86e1d62a109896e2c7e52c414e67d ]

    Terminate DMA transaction and clear CR3_DMAT when shutdown is requested,
    instead of when remove is requested. If DMA transfer is not stopped in
    shutdown ops, driver will fail to start a new DMA transfer after next
    startup ops.

    Fixes: 3489187204eb ("serial: stm32: adding dma support")
    Signed-off-by: Erwan Le Ray <erwan.leray@foss.st.com>
    Signed-off-by: Valentin Caron <valentin.caron@foss.st.com>
    Link: https://lore.kernel.org/r/20220104182445.4195-2-valentin.caron@foss.st.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4385d3b12808e086f75bdbd7642fa5f5d1ca01ea
Author: Alyssa Ross <hi@alyssa.is>
Date:   Tue Jan 4 13:10:28 2022 +0000

    serial: liteuart: fix MODULE_ALIAS

    [ Upstream commit 556172fabd226ba14b70c1740d0826a4717473dc ]

    modprobe can't handle spaces in aliases.

    Fixes: 1da81e5562fa ("drivers/tty/serial: add LiteUART driver")
    Signed-off-by: Alyssa Ross <hi@alyssa.is>
    Link: https://lore.kernel.org/r/20220104131030.1674733-1-hi@alyssa.is
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 124986477234ba3610cb9202d3c208939031b1bc
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Fri Dec 31 08:04:25 2021 +0000

    drivers/firmware: Add missing platform_device_put() in sysfb_create_simplefb

    [ Upstream commit 0589e8889dce8e0f0ea5bbf757f38865e2a469c1 ]

    Add the missing platform_device_put() before return from
    sysfb_create_simplefb() in the error handling case.

    Fixes: 8633ef82f101 ("drivers/firmware: consolidate EFI framebuffer setup for all
arches")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Link: https://lore.kernel.org/r/20211231080431.15385-1-linmq006@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e29581da866c650b4dce829d2694a92042cd1448
Author: Michal Suchanek <msuchanek@suse.de>
Date:   Tue Jan 4 18:05:05 2022 +0100

    debugfs: lockdown: Allow reading debugfs files that are not world readable

    [ Upstream commit 358fcf5ddbec4e6706405847d6a666f5933a6c25 ]

    When the kernel is locked down the kernel allows reading only debugfs
    files with mode 444. Mode 400 is also valid but is not allowed.

    Make the 444 into a mask.

    Fixes: 5496197f9b08 ("debugfs: Restrict debugfs when the kernel is locked down")
    Signed-off-by: Michal Suchanek <msuchanek@suse.de>
    Link: https://lore.kernel.org/r/20220104170505.10248-1-msuchanek@suse.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e4395ee7d07e37804a70712ba6306c129b6c4399
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Wed Jan 5 18:29:15 2022 +0100

    HID: hid-uclogic-params: Invalid parameter check in
uclogic_params_frame_init_v1_buttonpad

    [ Upstream commit aa320fdbbbb482c19100f51461bd0069753ce3d7 ]

    The function performs a check on the hdev input parameters, however, it
    is used before the check.

    Initialize the udev variable after the sanity check to avoid a
    possible NULL pointer dereference.

    Fixes: 9614219e9310e ("HID: uclogic: Extract tablet parameter discovery into a module")
    Addresses-Coverity-ID: 1443763 ("Null pointer dereference")
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f62b3be21ef50ac65defdf66239d9ce99282f8b7
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Wed Jan 5 18:29:14 2022 +0100

    HID: hid-uclogic-params: Invalid parameter check in uclogic_params_huion_init

    [ Upstream commit ff6b548afe4d9d1ff3a0f6ef79e8cbca25d8f905 ]
```

```
    The function performs a check on its input parameters, however, the
    hdev parameter is used before the check.

    Initialize the stack variables after checking the input parameters to
    avoid a possible NULL pointer dereference.

    Fixes: 9614219e9310e ("HID: uclogic: Extract tablet parameter discovery into a module")
    Addresses-Coverity-ID: 1443804 ("Null pointer dereference")
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e7e4dcace05228693a5bf53e94d833ce1cbce70
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Wed Jan 5 18:29:13 2022 +0100

    HID: hid-uclogic-params: Invalid parameter check in uclogic_params_get_str_desc

    [ Upstream commit 0a94131d6920916ccb6a357037c535533af08819 ]

    The function performs a check on the hdev input parameters, however, it
    is used before the check.

    Initialize the udev variable after the sanity check to avoid a
    possible NULL pointer dereference.

    Fixes: 9614219e9310e ("HID: uclogic: Extract tablet parameter discovery into a module")
    Addresses-Coverity-ID: 1443827 ("Null pointer dereference")
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e7de3630647585ee618dac55810ee25bd0eee813
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Wed Jan 5 18:29:12 2022 +0100

    HID: hid-uclogic-params: Invalid parameter check in uclogic_params_init

    [ Upstream commit f364c571a5c77e96de2d32062ff019d6b8d2e2bc ]

    The function performs a check on its input parameters, however, the
    hdev parameter is used before the check.

    Initialize the stack variables after checking the input parameters to
    avoid a possible NULL pointer dereference.

    Fixes: 9614219e9310e ("HID: uclogic: Extract tablet parameter discovery into a module")
    Addresses-Coverity-ID: 1443831 ("Null pointer dereference")
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92e6a5c95d5a5f0e79f26b452beb499728c65f95
Author: Pavel Hofman <pavel.hofman@ivitera.com>
Date:   Wed Jan 5 11:46:43 2022 +0100

    usb: gadget: u_audio: Subdevice 0 for capture ctls

    [ Upstream commit 601a5bc1aeef772ab1f47582fd322957799f5ab5 ]

    Both capture and playback alsa devices use subdevice 0. Yet capture-side
    ctls are defined for subdevice 1. The patch sets subdevice 0 for them.

    Fixes: 02de698ca812 ("usb: gadget: u_audio: add bi-directional volume and mute support")
```

```
    Signed-off-by: Pavel Hofman <pavel.hofman@ivitera.com>
    Link: https://lore.kernel.org/r/20220105104643.90125-1-pavel.hofman@ivitera.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a8e478e30466fd0444c6069c7e35cdcec4b8166
Author: John Keeping <john@metanate.com>
Date:   Tue Jan 4 18:32:42 2022 +0000

    usb: gadget: u_audio: fix calculations for small bInterval

    [ Upstream commit f2f69bf65df12176843ca11eab99949ba69e128b ]

    If bInterval is 1, then p_interval is 8000 and p_interval_mil is 8E9,
    which is too big for a 32-bit value.  While the storage is indeed
    64-bit, this value is used as the divisor in do_div() which will
    truncate it into a uint32_t leading to incorrect calculated values.

    Switch back to keeping the base value in struct snd_uac_chip which fits
    easily into an int, meaning that the division can be done in two steps
    with the divisor fitting safely into a uint32_t on both steps.

    Fixes: 6fec018a7e70 ("usb: gadget: u_audio.c: Adding Playback Pitch ctl for sync
playback")
    Tested-by: Pavel Hofman <pavel.hofman@ivitera.com>
    Signed-off-by: John Keeping <john@metanate.com>
    Link: https://lore.kernel.org/r/20220104183243.718258-1-john@metanate.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2055cc51bc4a25b378e888b39be19c606baf4eed
Author: John Keeping <john@metanate.com>
Date:   Thu Jan 6 11:57:31 2022 +0000

    usb: dwc2: gadget: initialize max_speed from params

    [ Upstream commit 92ef98a4caacad6d4a1490dda45d81ae5ccf5bc9 ]

    DWC2 may be paired with a full-speed PHY which is not capable of
    high-speed operation.  Report this correctly to the gadget core by
    setting max_speed from the core parameters.

    Prior to commit 5324bad66f09f ("usb: dwc2: gadget: implement
    udc_set_speed()") this didn't cause the hardware to be configured
    incorrectly, although the speed may have been reported incorrectly.  But
    after that commit params.speed is updated based on a value passed in by
    the gadget core which may set it to a faster speed than is supported by
    the hardware.  Initialising the max_speed parameter ensures the speed
    passed to dwc2_gadget_set_speed() will be one supported by the hardware.

    Fixes: 5324bad66f09f ("usb: dwc2: gadget: implement udc_set_speed()")
    Acked-by: Minas Harutyunyan <Minas.Harutyunyan@synopsys.com>
    Signed-off-by: John Keeping <john@metanate.com>
    Link: https://lore.kernel.org/r/20220106115731.1473909-1-john@metanate.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94aea60a03a62cbc54942b2e8fa933b8a4f2424b
Author: Dinh Nguyen <dinguyen@kernel.org>
Date:   Tue Jan 4 07:59:22 2022 -0600

    usb: dwc2: do not gate off the hardware if it does not support clock gating

    [ Upstream commit 34146c68083f1aef6709196b3dc888c1ceffd357 ]
```

```
    We should not be clearing the HCD_FLAG_HW_ACCESSIBLE bit if the hardware
    does not support clock gating.

    Fixes: 50fb0c128b6e ("usb: dwc2: Add clock gating entering flow by system suspend")
    Acked-by: Minas Harutyunyan <Minas.Harutyunyan@synopsys.com>
    Signed-off-by: Dinh Nguyen <dinguyen@kernel.org>
    Link: https://lore.kernel.org/r/20220104135922.734776-1-dinguyen@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d0ed1113ba26a515af47847cceb3618e8483595a
Author: Miaoqian Lin <linmq006@gmail.com>
Date:   Wed Dec 22 11:18:23 2021 +0000

    usb: dwc3: qcom: Fix NULL vs IS_ERR checking in dwc3_qcom_probe

    [ Upstream commit b52fe2dbb3e655eb1483000adfab68a219549e13 ]

    Since the acpi_create_platform_device() function may return error
    pointers, dwc3_qcom_create_urs_usb_platdev() function may return error
    pointers too. Using IS_ERR_OR_NULL() to check the return value to fix this.

    Fixes: c25c210f590e ("usb: dwc3: qcom: add URS Host support for sdm845 ACPI boot")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Link: https://lore.kernel.org/r/20211222111823.22887-1-linmq006@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac2e01b7e82a9f1f2185c86b2791e654a62b784e
Author: Wen Gu <guwen@linux.alibaba.com>
Date:   Thu Jan 6 20:42:08 2022 +0800

    net/smc: Reset conn->lgr when link group registration fails

    [ Upstream commit 36595d8ad46d9e4c41cc7c48c4405b7c3322deac ]

    SMC connections might fail to be registered in a link group due to
    unable to find a usable link during its creation. As a result,
    smc_conn_create() will return a failure and most resources related
    to the connection won't be applied or initialized, such as
    conn->abort_work or conn->lnk.

    If smc_conn_free() is invoked later, it will try to access the
    uninitialized resources related to the connection, thus causing
    a warning or crash.

    This patch tries to fix this by resetting conn->lgr to NULL if an
    abnormal exit occurs in smc_lgr_register_conn(), thus avoiding the
    access to uninitialized resources in smc_conn_free().

    Meanwhile, the new created link group should be terminated if smc
    connections can't be registered in it. So smc_lgr_cleanup_early() is
    modified to take care of link group only and invoked to terminate
    unusable link group by smc_conn_create(). The call to smc_conn_free()
    is moved out from smc_lgr_cleanup_early() to smc_conn_abort().

    Fixes: 56bc3b2094b4 ("net/smc: assign link to a new connection")
    Suggested-by: Karsten Graul <kgraul@linux.ibm.com>
    Signed-off-by: Wen Gu <guwen@linux.alibaba.com>
    Acked-by: Karsten Graul <kgraul@linux.ibm.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 85446a3b87799d87e6839611e5f528331bbe88fb
Author: Miaoqian Lin <linmq006@gmail.com>
```

```
Date:   Fri Dec 24 08:02:49 2021 +0000

    Bluetooth: hci_qca: Fix NULL vs IS_ERR_OR_NULL check in qca_serdev_probe

    [ Upstream commit 6845667146a28c09b5dfc401c1ad112374087944 ]

    The function devm_gpiod_get_index() return error pointers on error.
    Thus devm_gpiod_get_index_optional() could return NULL and error pointers.
    The same as devm_gpiod_get_optional() function. Using IS_ERR_OR_NULL()
    check to catch error pointers.

    Fixes: 77131dfe ("Bluetooth: hci_qca: Replace devm_gpiod_get() with
devm_gpiod_get_optional()")
    Signed-off-by: Miaoqian Lin <linmq006@gmail.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b627e93aced5ea2979176b5dce1a04cb3af42bb3
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Fri Dec 24 10:53:18 2021 +0800

    Bluetooth: hci_bcm: Check for error irq

    [ Upstream commit b38cd3b42fba66cc538edb9cf77e07881f43f8e2 ]

    For the possible failure of the platform_get_irq(), the returned irq
    could be error number and will finally cause the failure of the
    request_irq().
    Consider that platform_get_irq() can now in certain cases return
    -EPROBE_DEFER, and the consequences of letting request_irq() effectively
    convert that into -EINVAL, even at probe time rather than later on.
    So it might be better to check just now.

    Fixes: 0395ffc1ee05 ("Bluetooth: hci_bcm: Add PM for BCM devices")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d0afe10e0f2afc291ea63057b20934857adc4378
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Thu Jan 6 18:04:10 2022 +0800

    fsl/fman: Check for null pointer after calling devm_ioremap

    [ Upstream commit d5a73ec96cc57cf67e51b12820fc2354e7ca46f8 ]

    As the possible failure of the allocation, the devm_ioremap() may return
    NULL pointer.
    Take tgec_initialization() as an example.
    If allocation fails, the params->base_addr will be NULL pointer and will
    be assigned to tgec->regs in tgec_config().
    Then it will cause the dereference of NULL pointer in set_mac_address(),
    which is called by tgec_init().
    Therefore, it should be better to add the sanity check after the calling
    of the devm_ioremap().

    Fixes: 3933961682a3 ("fsl/fman: Add FMan MAC driver")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a52e0658d17a62ccfbe47d6a1d7db2d185446f3c
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Tue Jan 4 23:06:28 2022 +0800
```

```
    staging: greybus: audio: Check null pointer

    [ Upstream commit 2e81948177d769106754085c3e03534e6cc1f623 ]

    As the possible alloc failure of devm_kcalloc(), it could return null
    pointer.
    Therefore, 'strings' should be checked and return NULL if alloc fails to
    prevent the dereference of the NULL pointer.
    Also, the caller should also deal with the return value of the
    gb_generate_enum_strings() and return -ENOMEM if returns NULL.
    Moreover, because the memory allocated with devm_kzalloc() will be
    freed automatically when the last reference to the device is dropped,
    the 'gbe' in gbaudio_tplg_create_enum_kctl() and
    gbaudio_tplg_create_enum_ctl() do not need to free manually.
    But the 'control' in gbaudio_tplg_create_widget() and
    gbaudio_tplg_process_kcontrols() has a specially error handle to
    cleanup.
    So it should be better to cleanup 'control' when fails.

    Fixes: e65579e335da ("greybus: audio: topology: Enable enumerated control support")
    Reviewed-by: Alex Elder <elder@linaro.org>
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Link: https://lore.kernel.org/r/20220104150628.1987906-1-jiasheng@iscas.ac.cn
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e9a861cc0783cec4ffbbf550c96bc8d78e4943d
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Thu Jan 6 14:57:54 2022 +0300

    rocker: fix a sleeping in atomic bug

    [ Upstream commit 43d012123122cc69feacab55b71369f386c19566 ]

    This code is holding the &ofdpa->flow_tbl_lock spinlock so it is not
    allowed to sleep.  That means we have to pass the OFDPA_OP_FLAG_NOWAIT
    flag to ofdpa_flow_tbl_del().

    Fixes: 936bd486564a ("rocker: use FIB notifications instead of switchdev calls")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0089fa0b08d5c9fa88f9ec10ba79e59145feccf5
Author: Eric Dumazet <edumazet@google.com>
Date:   Wed Jan 5 03:48:42 2022 -0800

    ppp: ensure minimum packet size in ppp_write()

    [ Upstream commit 44073187990d5629804ce0627525f6ea5cfef171 ]

    It seems pretty clear ppp layer assumed user space
    would always be kind to provide enough data
    in their write() to a ppp device.

    This patch makes sure user provides at least
    2 bytes.

    It adds PPP_PROTO_LEN macro that could replace
    in net-next many occurrences of hard-coded 2 value.

    I replaced only one occurrence to ease backports
    to stable kernels.

    The bug manifests in the following report:
```

```
    BUG: KMSAN: uninit-value in ppp_send_frame+0x28d/0x27c0
drivers/net/ppp/ppp_generic.c:1740
     ppp_send_frame+0x28d/0x27c0 drivers/net/ppp/ppp_generic.c:1740
     __ppp_xmit_process+0x23e/0x4b0 drivers/net/ppp/ppp_generic.c:1640
     ppp_xmit_process+0x1fe/0x480 drivers/net/ppp/ppp_generic.c:1661
     ppp_write+0x5cb/0x5e0 drivers/net/ppp/ppp_generic.c:513
     do_iter_write+0xb0c/0x1500 fs/read_write.c:853
     vfs_writev fs/read_write.c:924 [inline]
     do_writev+0x645/0xe00 fs/read_write.c:967
     __do_sys_writev fs/read_write.c:1040 [inline]
     __se_sys_writev fs/read_write.c:1037 [inline]
     __x64_sys_writev+0xe5/0x120 fs/read_write.c:1037
     do_syscall_x64 arch/x86/entry/common.c:51 [inline]
     do_syscall_64+0x54/0xd0 arch/x86/entry/common.c:82
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    Uninit was created at:
     slab_post_alloc_hook mm/slab.h:524 [inline]
     slab_alloc_node mm/slub.c:3251 [inline]
     __kmalloc_node_track_caller+0xe0c/0x1510 mm/slub.c:4974
     kmalloc_reserve net/core/skbuff.c:354 [inline]
     __alloc_skb+0x545/0xf90 net/core/skbuff.c:426
     alloc_skb include/linux/skbuff.h:1126 [inline]
     ppp_write+0x11d/0x5e0 drivers/net/ppp/ppp_generic.c:501
     do_iter_write+0xb0c/0x1500 fs/read_write.c:853
     vfs_writev fs/read_write.c:924 [inline]
     do_writev+0x645/0xe00 fs/read_write.c:967
     __do_sys_writev fs/read_write.c:1040 [inline]
     __se_sys_writev fs/read_write.c:1037 [inline]
     __x64_sys_writev+0xe5/0x120 fs/read_write.c:1037
     do_syscall_x64 arch/x86/entry/common.c:51 [inline]
     do_syscall_64+0x54/0xd0 arch/x86/entry/common.c:82
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Paul Mackerras <paulus@samba.org>
    Cc: linux-ppp@vger.kernel.org
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Acked-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 675bc9e8329de0f1c76d0b66358b1ba2cb5c6afd
Author: Miroslav Lichvar <mlichvar@redhat.com>
Date:   Wed Jan 5 11:33:26 2022 +0100

    net: fix SOF_TIMESTAMPING_BIND_PHC to work with multiple sockets

    [ Upstream commit 007747a984ea5e895b7d8b056b24ebf431e1e71d ]

    When multiple sockets using the SOF_TIMESTAMPING_BIND_PHC flag received
    a packet with a hardware timestamp (e.g. multiple PTP instances in
    different PTP domains using the UDPv4/v6 multicast or L2 transport),
    the timestamps received on some sockets were corrupted due to repeated
    conversion of the same timestamp (by the same or different vclocks).

    Fix ptp_convert_timestamp() to not modify the shared skb timestamp
    and return the converted timestamp as a ktime_t instead. If the
    conversion fails, return 0 to not confuse the application with
    timestamps corresponding to an unexpected PHC.

    Fixes: d7c088265588 ("net: socket: support hardware timestamp conversion to PHC bound")
    Signed-off-by: Miroslav Lichvar <mlichvar@redhat.com>
```

```
    Cc: Yangbo Lu <yangbo.lu@nxp.com>
    Cc: Richard Cochran <richardcochran@gmail.com>
    Acked-by: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37b715c8f483a15280d40adfc30c287301bf3310
Author: Florian Westphal <fw@strlen.de>
Date:   Wed Jan 5 14:19:54 2022 +0100

    netfilter: nft_set_pipapo: allocate pcpu scratch maps on clone

    [ Upstream commit 23c54263efd7cb605e2f7af72717a2a951999217 ]

    This is needed in case a new transaction is made that doesn't insert any
    new elements into an already existing set.

    Else, after second 'nft -f ruleset.txt', lookups in such a set will fail
    because ->lookup() encounters raw_cpu_ptr(m->scratch) == NULL.

    For the initial rule load, insertion of elements takes care of the
    allocation, but for rule reloads this isn't guaranteed: we might not
    have additions to the set.

    Fixes: 3c4287f62044a90e ("nf_tables: Add set type for arbitrary concatenation of
ranges")
    Reported-by: etkaar <lists.netfilter.org@prvy.eu>
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Reviewed-by: Stefano Brivio <sbrivio@redhat.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29daf52cd8cfb4c68cc9f27aa0001d94a4b9e432
Author: Pablo Neira Ayuso <pablo@netfilter.org>
Date:   Wed Jan 5 16:09:57 2022 +0100

    netfilter: nft_payload: do not update layer 4 checksum when mangling fragments

    [ Upstream commit 4e1860a3863707e8177329c006d10f9e37e097a8 ]

    IP fragments do not come with the transport header, hence skip bogus
    layer 4 checksum updates.

    Fixes: 1814096980bb ("netfilter: nft_payload: layer 4 checksum adjustment for
pseudoheader fields")
    Reported-and-tested-by: Steffen Weinreich <steve@weinreich.org>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8feb22063dca1874713d53b1549eca73b7f5e512
Author: Qiang Wang <wangqiang.wq.frank@bytedance.com>
Date:   Mon Dec 27 21:07:12 2021 +0800

    libbpf: Use probe_name for legacy kprobe

    [ Upstream commit 71cff670baff5cc6a6eeb0181e2cc55579c5e1e0 ]

    Fix a bug in commit 46ed5fc33db9, which wrongly used the
    func_name instead of probe_name to register legacy kprobe.

    Fixes: 46ed5fc33db9 ("libbpf: Refactor and simplify legacy kprobe code")
    Co-developed-by: Chengming Zhou <zhouchengming@bytedance.com>
    Signed-off-by: Qiang Wang <wangqiang.wq.frank@bytedance.com>
    Signed-off-by: Chengming Zhou <zhouchengming@bytedance.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
```

```
    Tested-by: Hengqi Chen <hengqi.chen@gmail.com>
    Reviewed-by: Hengqi Chen <hengqi.chen@gmail.com>
    Link: https://lore.kernel.org/bpf/20211227130713.66933-1-
 wangqiang.wq.frank@bytedance.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 57486171877c6c7bb4a322f5f760155b2778fd0e
 Author: Kuniyuki Iwashima <kuniyu@amazon.co.jp>
 Date:   Tue Jan 4 10:31:48 2022 +0900

    bpf: Fix SO_RCVBUF/SO_SNDBUF handling in _bpf_setsockopt().

    [ Upstream commit 04c350b1ae6bdb12b84009a4d0bf5ab4e621c47b ]

    The commit 4057765f2dee ("sock: consistent handling of extreme
    SO_SNDBUF/SO_RCVBUF values") added a change to prevent underflow
    in setsockopt() around SO_SNDBUF/SO_RCVBUF.

    This patch adds the same change to _bpf_setsockopt().

    Fixes: 4057765f2dee ("sock: consistent handling of extreme SO_SNDBUF/SO_RCVBUF values")
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.co.jp>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Link: https://lore.kernel.org/bpf/20220104013153.97906-2-kuniyu@amazon.co.jp
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 6955b4139f1ca59ccdbdea6081e9ab28df47cd28
 Author: Kris Van Hees <kris.van.hees@oracle.com>
 Date:   Wed Jan 5 16:01:50 2022 -0500

    bpf: Fix verifier support for validation of async callbacks

    [ Upstream commit a5bebc4f00dee47113eed48098c68e88b5ba70e8 ]

    Commit bfc6bb74e4f1 ("bpf: Implement verifier support for validation of async
 callbacks.")
    added support for BPF_FUNC_timer_set_callback to
    the __check_func_call() function.  The test in __check_func_call() is
    flaweed because it can mis-interpret a regular BPF-to-BPF pseudo-call
    as a BPF_FUNC_timer_set_callback callback call.

    Consider the conditional in the code:

            if (insn->code == (BPF_JMP | BPF_CALL) &&
                insn->imm == BPF_FUNC_timer_set_callback) {

    The BPF_FUNC_timer_set_callback has value 170.  This means that if you
    have a BPF program that contains a pseudo-call with an instruction delta
    of 170, this conditional will be found to be true by the verifier, and
    it will interpret the pseudo-call as a callback.  This leads to a mess
    with the verification of the program because it makes the wrong
    assumptions about the nature of this call.

    Solution: include an explicit check to ensure that insn->src_reg == 0.
    This ensures that calls cannot be mis-interpreted as an async callback
    call.

    Fixes: bfc6bb74e4f1 ("bpf: Implement verifier support for validation of async
 callbacks.")
    Signed-off-by: Kris Van Hees <kris.van.hees@oracle.com>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Link: https://lore.kernel.org/bpf/20220105210150.GH1559@oracle.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 2bdeb53e6a1f8addf6be9deb868d32d80b51c671
```

```
Author: Daniel Borkmann <daniel@iogearbox.net>
Date:   Wed Jan 5 11:35:13 2022 -0800

    bpf: Don't promote bogus looking registers after null check.

    [ Upstream commit e60b0d12a95dcf16a63225cead4541567f5cb517 ]

    If we ever get to a point again where we convert a bogus looking <ptr>_or_null
    typed register containing a non-zero fixed or variable offset, then lets not
    reset these bounds to zero since they are not and also don't promote the register
    to a <ptr> type, but instead leave it as <ptr>_or_null. Converting to a unknown
    register could be an avenue as well, but then if we run into this case it would
    allow to leak a kernel pointer this way.

    Fixes: f1174f77b50c ("bpf/verifier: rework value tracking")
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88d8fdad8259590c4bfe66ada7d5c6526ddb295e
Author: John Fastabend <john.fastabend@gmail.com>
Date:   Tue Jan 4 13:46:45 2022 -0800

    bpf, sockmap: Fix double bpf_prog_put on error case in map_link

    [ Upstream commit 218d747a4142f281a256687bb513a135c905867b ]

    sock_map_link() is called to update a sockmap entry with a sk. But, if the
    sock_map_init_proto() call fails then we return an error to the map_update
    op against the sockmap. In the error path though we need to cleanup psock
    and dec the refcnt on any programs associated with the map, because we
    refcnt them early in the update process to ensure they are pinned for the
    psock. (This avoids a race where user deletes programs while also updating
    the map with new socks.)

    In current code we do the prog refcnt dec explicitely by calling
    bpf_prog_put() when the program was found in the map. But, after commit
    '38207a5e81230' in this error path we've already done the prog to psock
    assignment so the programs have a reference from the psock as well. This
    then causes the psock tear down logic, invoked by sk_psock_put() in the
    error path, to similarly call bpf_prog_put on the programs there.

    To be explicit this logic does the prog->psock assignment:

      if (msg_*)
        psock_set_prog(...)

    Then the error path under the out_progs label does a similar check and
    dec with:

      if (msg_*)
         bpf_prog_put(...)

    And the teardown logic sk_psock_put() does ...

      psock_set_prog(msg_*, NULL)

    ... triggering another bpf_prog_put(...). Then KASAN gives us this splat,
    found by syzbot because we've created an inbalance between bpf_prog_inc and
    bpf_prog_put calling put twice on the program.

      BUG: KASAN: vmalloc-out-of-bounds in __bpf_prog_put kernel/bpf/syscall.c:1812 [inline]
      BUG: KASAN: vmalloc-out-of-bounds in __bpf_prog_put kernel/bpf/syscall.c:1812 [inline]
kernel/bpf/syscall.c:1829
      BUG: KASAN: vmalloc-out-of-bounds in bpf_prog_put+0x8c/0x4f0 kernel/bpf/syscall.c:1829
```

```
kernel/bpf/syscall.c:1829
      Read of size 8 at addr ffffc90000e76038 by task syz-executor020/3641

    To fix clean up error path so it doesn't try to do the bpf_prog_put in the
    error path once progs are assigned then it relies on the normal psock
    tear down logic to do complete cleanup.

    For completness we also cover the case whereh sk_psock_init_strp() fails,
    but this is not expected because it indicates an incorrect socket type
    and should be caught earlier.

    Fixes: 38207a5e8123 ("bpf, sockmap: Attach map progs to psock early for feature probes")
    Reported-by: syzbot+bb73e71cf4b8fd376a4f@syzkaller.appspotmail.com
    Signed-off-by: John Fastabend <john.fastabend@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20220104214645.290900-1-john.fastabend@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19bba036f6a42e866d6ea48fc729113c46d4562d
Author: John Fastabend <john.fastabend@gmail.com>
Date:   Tue Jan 4 12:59:18 2022 -0800

    bpf, sockmap: Fix return codes from tcp_bpf_recvmsg_parser()

    [ Upstream commit 5b2c5540b8110eea0d67a78fb0ddb9654c58daeb ]

    Applications can be confused slightly because we do not always return the
    same error code as expected, e.g. what the TCP stack normally returns. For
    example on a sock err sk->sk_err instead of returning the sock_error we
    return EAGAIN. This usually means the application will 'try again'
    instead of aborting immediately. Another example, when a shutdown event
    is received we should immediately abort instead of waiting for data when
    the user provides a timeout.

    These tend to not be fatal, applications usually recover, but introduces
    bogus errors to the user or introduces unexpected latency. Before
    'c5d2177a72a16' we fell back to the TCP stack when no data was available
    so we managed to catch many of the cases here, although with the extra
    latency cost of calling tcp_msg_wait_data() first.

    To fix lets duplicate the error handling in TCP stack into tcp_bpf so
    that we get the same error codes.

    These were found in our CI tests that run applications against sockmap
    and do longer lived testing, at least compared to test_sockmap that
    does short-lived ping/pong tests, and in some of our test clusters
    we deploy.

    Its non-trivial to do these in a shorter form CI tests that would be
    appropriate for BPF selftests, but we are looking into it so we can
    ensure this keeps working going forward. As a preview one idea is to
    pull in the packetdrill testing which catches some of this.

    Fixes: c5d2177a72a16 ("bpf, sockmap: Fix race in ingress receive verdict with redirect
to self")
    Signed-off-by: John Fastabend <john.fastabend@gmail.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20220104205918.286416-1-john.fastabend@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit af685851087855e446406eb5560c3331ea2c8b2e
Author: Hou Tao <houtao1@huawei.com>
Date:   Fri Dec 31 23:10:18 2021 +0800

    bpf, arm64: Use emit_addr_mov_i64() for BPF_PSEUDO_FUNC
```

```
    [ Upstream commit e4a41c2c1fa916547e63440c73a51a5eb06247af ]

    The following error is reported when running "./test_progs -t for_each"
    under arm64:

      bpf_jit: multi-func JIT bug 58 != 56
      [...]
      JIT doesn't support bpf-to-bpf calls

    The root cause is the size of BPF_PSEUDO_FUNC instruction increases
    from 2 to 3 after the address of called bpf-function is settled and
    there are two bpf-to-bpf calls in test_pkt_access. The generated
    instructions are shown below:

      0x48:  21 00 C0 D2    movz x1, #0x1, lsl #32
      0x4c:  21 00 80 F2    movk x1, #0x1

      0x48:  E1 3F C0 92    movn x1, #0x1ff, lsl #32
      0x4c:  41 FE A2 F2    movk x1, #0x17f2, lsl #16
      0x50:  81 70 9F F2    movk x1, #0xfb84

    Fixing it by using emit_addr_mov_i64() for BPF_PSEUDO_FUNC, so
    the size of jited image will not change.

    Fixes: 69c087ba6225 ("bpf: Add bpf_for_each_map_elem() helper")
    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211231151018.3781550-1-houtao1@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 26543e3dead794a29686a44118a699ad6f3d91f4
Author: Xin Xiong <xiongx18@fudan.edu.cn>
Date:   Thu Dec 23 10:48:12 2021 +0800

    netfilter: ipt_CLUSTERIP: fix refcount leak in clusterip_tg_check()

    [ Upstream commit d94a69cb2cfa77294921aae9afcfb866e723a2da ]

    The issue takes place in one error path of clusterip_tg_check(). When
    memcmp() returns nonzero, the function simply returns the error code,
    forgetting to decrease the reference count of a clusterip_config
    object, which is bumped earlier by clusterip_config_find_get(). This
    may incur reference count leak.

    Fix this issue by decrementing the refcount of the object in specific
    error path.

    Fixes: 06aa151ad1fc74 ("netfilter: ipt_CLUSTERIP: check MAC address when duplicate
config is set")
    Signed-off-by: Xin Xiong <xiongx18@fudan.edu.cn>
    Signed-off-by: Xiyu Yang <xiyuyang19@fudan.edu.cn>
    Signed-off-by: Xin Tan <tanxin.ctf@gmail.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7e84bfcf949cf8e1b6d5b4df76441f055a15c6e
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Wed Jan 5 15:18:11 2022 +0200

    net: dsa: fix incorrect function pointer check for MRP ring roles

    [ Upstream commit ff91e1b68490b97c18c649b769618815eb945f11 ]

    The cross-chip notifier boilerplate code meant to check the presence of
```

```
    ds->ops->port_mrp_add_ring_role before calling it, but checked
    ds->ops->port_mrp_add instead, before calling
    ds->ops->port_mrp_add_ring_role.

    Therefore, a driver which implements one operation but not the other
    would trigger a NULL pointer dereference.

    There isn't any such driver in DSA yet, so there is no reason to
    backport the change. Issue found through code inspection.

    Cc: Horatiu Vultur <horatiu.vultur@microchip.com>
    Fixes: c595c4330da0 ("net: dsa: add MRP support")
    Signed-off-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 4760f77517a57767b6519c0c87da02d6c51bcead
Author: Daniel Golle <daniel@makrotopia.org>
Date:   Tue Jan 4 12:06:22 2022 +0000

    net: ethernet: mtk_eth_soc: fix return values and refactor MDIO ops

    [ Upstream commit eda80b249df7bbc7b3dd13907343a3e59bfc57fd ]

    Instead of returning -1 (-EPERM) when MDIO bus is stuck busy
    while writing or 0xffff if it happens while reading, return the
    appropriate -ETIMEDOUT. Also fix return type to int instead of u32.
    Refactor functions to use bitfield helpers instead of having various
    masking and shifting constants in the code, which also results in the
    register definitions in the header file being more obviously related
    to what is stated in the MediaTek's Reference Manual.

    Fixes: 656e705243fd0 ("net-next: mediatek: add support for MT7623 ethernet")
    Signed-off-by: Daniel Golle <daniel@makrotopia.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a2ae51a58e6ad7bc19212bb90c636ad7cddabba
Author: Raed Salem <raeds@nvidia.com>
Date:   Mon Jan 3 13:19:29 2022 +0200

    net/xfrm: IPsec tunnel mode fix inner_ipproto setting in sec_path

    [ Upstream commit 45a98ef4922def8c679ca7c454403d1957fe70e7 ]

    The inner_ipproto saves the inner IP protocol of the plain
    text packet. This allows vendor's IPsec feature making offload
    decision at skb's features_check and configuring hardware at
    ndo_start_xmit, current code implenetation did not handle the
    case where IPsec is used in tunnel mode.

    Fix by handling the case when IPsec is used in tunnel mode by
    reading the protocol of the plain text packet IP protocol.

    Fixes: fa4535238fb5 ("net/xfrm: Add inner_ipproto into sec_path")
    Signed-off-by: Raed Salem <raeds@nvidia.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58d7756d14734993cf80cb7b68781f0dfe855765
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Mon Dec 20 16:38:11 2021 +0800

    power: reset: mt6397: Check for null res pointer
```

```
    [ Upstream commit 1c1348bf056dee665760a3bd1cd30b0be7554fc2 ]

    The return value of platform_get_resource() needs to be checked.
    To avoid use of error pointer in case that there is no suitable
    resource.

    Fixes: d28c74c10751 ("power: reset: add driver for mt6323 poweroff")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1b5c8bfdccf86502abc8db5809f636ec9b18c8e
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Wed Dec 1 02:11:40 2021 +0800

    pcmcia: rsrc_nonstatic: Fix a NULL pointer dereference in nonstatic_find_mem_region()

    [ Upstream commit 977d2e7c63c3d04d07ba340b39987742e3241554 ]

    In nonstatic_find_mem_region(), pcmcia_make_resource() is assigned to
    res and used in pci_bus_alloc_resource(). There a dereference of res
    in pci_bus_alloc_resource(), which could lead to a NULL pointer
    dereference on failure of pcmcia_make_resource().

    Fix this bug by adding a check of res.

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_PCCARD_NONSTATIC=y show no new warnings,
    and our static analyzer no longer warns about this code.

    Fixes: 49b1153adfe1 ("pcmcia: move all pcmcia_resource_ops providers into one module")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Dominik Brodowski <linux@dominikbrodowski.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 99dcb51be8e60dbcd059abf55eec3865398a6b73
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Wed Dec 1 00:59:23 2021 +0800

    pcmcia: rsrc_nonstatic: Fix a NULL pointer dereference in __nonstatic_find_io_region()

    [ Upstream commit ca0fe0d7c35c97528bdf621fdca75f13157c27af ]

    In __nonstatic_find_io_region(), pcmcia_make_resource() is assigned to
    res and used in pci_bus_alloc_resource(). There is a dereference of res
    in pci_bus_alloc_resource(), which could lead to a NULL pointer
    dereference on failure of pcmcia_make_resource().

    Fix this bug by adding a check of res.

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.
```

```
    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_PCCARD_NONSTATIC=y show no new warnings,
    and our static analyzer no longer warns about this code.

    Fixes: 49b1153adfe1 ("pcmcia: move all pcmcia_resource_ops providers into one module")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    [linux@dominikbrodowski.net: Fix typo in commit message]
    Signed-off-by: Dominik Brodowski <linux@dominikbrodowski.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf8272b41ff357c085b0cf3cc007145a1236ad37
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Dec 30 12:57:47 2021 +0100

    ACPI: scan: Create platform device for BCM4752 and LNV4752 ACPI nodes

    [ Upstream commit f85196bdd5a50da74670250564740fc852b3c239 ]

    BCM4752 and LNV4752 ACPI nodes describe a Broadcom 4752 GPS module
    attached to an UART of the system.

    The GPS modules talk a custom protocol which only works with a closed-
    source Android gpsd daemon which knows this protocol.

    The ACPI nodes also describe GPIOs to turn the GPS on/off these are
    handled by the net/rfkill/rfkill-gpio.c code. This handling predates the
    addition of enumeration of ACPI instantiated serdevs to the kernel and
    was broken by that addition, because the ACPI scan code now no longer
    instantiates platform_device-s for these nodes.

    Rename the i2c_multi_instantiate_ids HID list to ignore_serial_bus_ids
    and add the BCM4752 and LNV4752 HIDs, so that rfkill-gpio gets
    a platform_device to bind to again; and so that a tty cdev for gpsd
    gets created for these.

    Fixes: e361d1f85855 ("ACPI / scan: Fix enumeration for special UART devices")
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59728d43f1b0a3129318647640e1232591f52522
Author: José Expósito <jose.exposito89@gmail.com>
Date:   Sun Dec 26 12:16:14 2021 +0100

    drm/amd/display: fix dereference before NULL check

    [ Upstream commit f28cad86ada1a7345d7bbd379bef5a8babfa791b ]

    The "plane_state" pointer was access before checking if it was NULL.

    Avoid a possible NULL pointer dereference by accessing the plane
    address after the check.

    Addresses-Coverity-ID: 1493892 ("Dereference before null check")
    Fixes: 3f68c01be9a22 ("drm/amd/display: add cyan_skillfish display support")
    Signed-off-by: José Expósito <jose.exposito89@gmail.com>
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50c66ef7d30bc41f41e30c72d8133be8c0ee8929
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
```

```
Date:   Fri Dec 24 14:29:09 2021 +0000

    serial: 8250_bcm7271: Propagate error codes from brcmuart_probe()

    [ Upstream commit c195438f1e84de8fa46b4f5264d12379bee6e9a1 ]

    In case of failures brcmuart_probe() always returned -ENODEV, this
    isn't correct for example platform_get_irq_byname() may return
    -EPROBE_DEFER to handle such cases propagate error codes in
    brcmuart_probe() in case of failures.

    Fixes: 41a469482de25 ("serial: 8250: Add new 8250-core based Broadcom STB driver")
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Link: https://lore.kernel.org/r/20211224142917.6966-4-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f7a6dd58e0817b063252d7c5bec88e588df34b31
Author: Haimin Zhang <tcs_kernel@tencent.com>
Date:   Wed Dec 29 19:20:02 2021 +0800

    bpf: Add missing map_get_next_key method to bloom filter map.

    [ Upstream commit 3ccdcee28415c4226de05438b4d89eb5514edf73 ]

    Without it, kernel crashes in map_get_next_key().

    Fixes: 9330986c0300 ("bpf: Add bloom filter map implementation")
    Reported-by: TCS Robot <tcs_robot@tencent.com>
    Signed-off-by: Haimin Zhang <tcs_kernel@tencent.com>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Joanne Koong <joannekoong@fb.com>
    Link: https://lore.kernel.org/bpf/1640776802-22421-1-git-send-email-tcs.kernel@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3bdc5e1997c916ccc87d53858e016cab1a7ec2f
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Wed Dec 15 22:08:44 2021 +0000

    io_uring: remove double poll on poll update

    [ Upstream commit e840b4baf3cfb37e2ead4f649a45bb78178677ff ]

    Before updating a poll request we should remove it from poll queues,
    including the double poll entry.

    Fixes: b69de288e913 ("io_uring: allow events and user_data update of running poll
requests")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/ac39e7f80152613603b8a6cc29a2b6063ac2434f.1639605189.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 915b4f3e22cdc96c72953af73cbdfca0b2ed14f1
Author: Zhang Zixun <zhang133010@icloud.com>
Date:   Mon Dec 27 22:02:49 2021 +0100

    x86/mce/inject: Avoid out-of-bounds write when setting flags

    [ Upstream commit de768416b203ac84e02a757b782a32efb388476f ]
```

```
    A contrived zero-length write, for example, by using write(2):

      ...
      ret = write(fd, str, 0);
      ...

    to the "flags" file causes:

      BUG: KASAN: stack-out-of-bounds in flags_write
      Write of size 1 at addr ffff888019be7ddf by task writefile/3787

      CPU: 4 PID: 3787 Comm: writefile Not tainted 5.16.0-rc7+ #12
      Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.14.0-2 04/01/2014

    due to accessing buf one char before its start.

    Prevent such out-of-bounds access.

      [ bp: Productize into a proper patch. Link below is the next best
        thing because the original mail didn't get archived on lore. ]

    Fixes: 0451d14d0561 ("EDAC, mce_amd_inj: Modify flags attribute to use string
arguments")
    Signed-off-by: Zhang Zixun <zhang133010@icloud.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/linux-edac/YcnePfF1OOqoQwrX@zn.tnic/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6857d1742fee225a9a0389ea6d6343ab94d820d6
Author: Arseny Demidov <arsdemal@gmail.com>
Date:   Sun Dec 19 13:22:39 2021 +0300

    hwmon: (mr75203) fix wrong power-up delay value

    [ Upstream commit a8d6d4992ad9d92356619ac372906bd29687bb46 ]

    In the file mr75203.c we have a macro named POWER_DELAY_CYCLE_256,
    the correct value should be 0x100. The register ip_tmr is expressed
    in units of IP clk cycles, in accordance with the datasheet.
    Typical power-up delays for Temperature Sensor are 256 cycles i.e. 0x100.

    Fixes: 9d823351a337 ("hwmon: Add hardware monitoring driver for Moortec MR75203 PVT
controller")
    Signed-off-by: Arseny Demidov <a.demidov@yadro.com>
    Link: https://lore.kernel.org/r/20211219102239.1112-1-a.demidov@yadro.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 81c6534ce9b0a526cec5b790c8b0718a077ea4d6
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Fri Dec 24 12:34:50 2021 +0100

    regulator: qcom-labibb: OCP interrupts are not a failure while disabled

    [ Upstream commit d27bb69dc83f00f86a830298c967052cded6e784 ]

    Receiving the Over-Current Protection interrupt while the regulator is
    disabled does not count as unhandled/failure (IRQ_NONE, or 0 as it were)
    but a "fake event", usually due to inrush as the is regulator about to
    be enabled.

    Fixes: 390af53e0411 ("regulator: qcom-labibb: Implement short-circuit and over-current
IRQs")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
```

```
    Link: https://lore.kernel.org/r/20211224113450.107958-1-marijn.suijten@somainline.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d5d6e92af32d20582a429f87f0aeb46d031e20b1
Author: Ming Lei <ming.lei@redhat.com>
Date:   Fri Dec 24 09:08:31 2021 +0800

    block: null_blk: only set set->nr_maps as 3 if active poll_queues is > 0

    [ Upstream commit 19768f80cf23834e65482f1667ff54192d469fee ]

    It isn't correct to set set->nr_maps as 3 if g_poll_queues is > 0 since
    we can change it via configfs for null_blk device created there, so only
    set it as 3 if active poll_queues is > 0.

    Fixes divide zero exception reported by Shinichiro.

    Fixes: 2bfdbe8b7ebd ("null_blk: allow zero poll queues")
    Reported-by: Shinichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Reviewed-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Link: https://lore.kernel.org/r/20211224010831.1521805-1-ming.lei@redhat.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3475ba69723575efb4b2ed0bb7157deb6485d0c9
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Fri Dec 17 10:13:32 2021 +0300

    crypto: octeontx2 - prevent underflow in get_cores_bmap()

    [ Upstream commit 10371b6212bb682f13247733d6b76b91b2b80f9a ]

    If we're going to cap "eng_grp->g->engs_num" upper bounds then we should
    cap the lower bounds as well.

    Fixes: 43ac0b824f1c ("crypto: octeontx2 - load microcode and create engine groups")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 52e5c80b049b56b698012ac50c0bb16498cbca3e
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed Dec 22 09:30:41 2021 -0700

    x86/boot/compressed: Move CLANG_FLAGS to beginning of KBUILD_CFLAGS

    [ Upstream commit 5fe392ff9d1f7254a1fbb3f72d9893088e4d23eb ]

    When cross compiling i386_defconfig on an arm64 host with clang, there
    are a few instances of '-Waddress-of-packed-member' and
    '-Wgnu-variable-sized-type-not-at-end' in arch/x86/boot/compressed/,
    which should both be disabled with the cc-disable-warning calls in that
    directory's Makefile, which indicates that cc-disable-warning is failing
    at the point of testing these flags.

    The cc-disable-warning calls fail because at the point that the flags
    are tested, KBUILD_CFLAGS has '-march=i386' without $(CLANG_FLAGS),
    which has the '--target=' flag to tell clang what architecture it is
    targeting. Without the '--target=' flag, the host architecture (arm64)
    is used and i386 is not a valid value for '-march=' in that case. This
    error can be seen by adding some logging to try-run:

      clang-14: error: the clang compiler does not support '-march=i386'
```

```
    Invoking the compiler has to succeed prior to calling cc-option or
    cc-disable-warning in order to accurately test whether or not the flag
    is supported; if it doesn't, the requested flag can never be added to
    the compiler flags. Move $(CLANG_FLAGS) to the beginning of KBUILD_FLAGS
    so that any new flags that might be added in the future can be
    accurately tested.

    Fixes: d5cbd80e302d ("x86/boot: Add $(CLANG_FLAGS) to compressed KBUILD_CFLAGS")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lore.kernel.org/r/20211222163040.1961481-1-nathan@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4618a4370dc016fe7b6dbbeb79edfa81779871b4
Author: Panicker Harish <quic_pharish@quicinc.com>
Date:   Wed Dec 22 12:59:05 2021 +0530

    Bluetooth: hci_qca: Stop IBS timer during BT OFF

    [ Upstream commit df1e5c51492fd93ffc293acdcc6f00698d19fedc ]

    The IBS timers are not stopped properly once BT OFF is triggered.
    we could see IBS commands being sent along with version command,
    so stopped IBS timers while Bluetooth is off.

    Fixes: 3e4be65eb82c ("Bluetooth: hci_qca: Add poweroff support during hci down for
wcn3990")
    Signed-off-by: Panicker Harish <quic_pharish@quicinc.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4eb478120ce92eb9e0e605478e67d22274c29545
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Tue Dec 21 14:57:09 2021 +0100

    mt76: mt7921: fix a possible race enabling/disabling runtime-pm

    [ Upstream commit d430dffbe9dd30759f3c64b65bf85b0245c8d8ab ]

    Fix a possible race enabling/disabling runtime-pm between
    mt7921_pm_set() and mt7921_poll_rx() since mt7921_pm_wake_work()
    always schedules rx-napi callback and it will trigger
    mt7921_pm_power_save_work routine putting chip to in low-power state
    during mt7921_pm_set processing.

    Suggested-by: Deren Wu <deren.wu@mediatek.com>
    Tested-by: Deren Wu <deren.wu@mediatek.com>
    Fixes: 1d8efc741df8 ("mt76: mt7921: introduce Runtime PM support")
    Signed-off-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link:
https://lore.kernel.org/r/0f3e075a2033dc05f09dab4059e5be8cbdccc239.1640094847.git.lorenzo@ke
rnel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1c6e91226867da8eb3cbc39073d99edc1e31893c
Author: Baochen Qiang <quic_bqiang@quicinc.com>
Date:   Wed Dec 22 09:35:35 2021 +0800

    ath11k: Fix unexpected return buffer manager error for QCA6390

    [ Upstream commit 71c748b5e01e3e28838a8e26a8966fb5adb03df7 ]

    We are seeing below error on QCA6390:
```

```
    ...
    [70211.671189] ath11k_pci 0000:72:00.0: failed to parse rx error in wbm_rel ring desc
-22
    [70212.696154] ath11k_pci 0000:72:00.0: failed to parse rx error in wbm_rel ring desc
-22
    [70213.092941] ath11k_pci 0000:72:00.0: failed to parse rx error in wbm_rel ring desc
-22
    ...

    The reason is that, with commit 734223d78428 ("ath11k: change return
    buffer manager for QCA6390"), ath11k expects the return buffer manager
    (RBM) field of descriptor configured as HAL_RX_BUF_RBM_SW1_BM when
    parsing error frames from WBM2SW3_RELEASE ring. This is a wrong change
    cause the RBM field is set as HAL_RX_BUF_RBM_SW3_BM.

    The same issue also applies to REO2TCL ring though we have not got any
    error reported.

    Fix it by changing RBM from HAL_RX_BUF_RBM_SW1_BM to HAL_RX_BUF_RBM_SW3_BM
    for these two rings.

    Tested-on: QCA6390 hw2.0 PCI WLAN.HST.1.0.1-01740-QCAHSTSWPLZ_V2_TO_X86-1

    Fixes: 734223d78428 ("ath11k: change return buffer manager for QCA6390")
    Signed-off-by: Baochen Qiang <quic_bqiang@quicinc.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20211222013536.582527-1-quic_bqiang@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db7b553cb1afa685b164e9a1f94f3e99c78312a2
Author: Clément Léger <clement.leger@bootlin.com>
Date:   Mon Dec 20 22:05:33 2021 +0100

    software node: fix wrong node passed to find nargs_prop

    [ Upstream commit c5fc5ba8b6b7bebc05e45036a33405b4c5036c2f ]

    nargs_prop refers to a property located in the reference that is found
    within the nargs property. Use the correct reference node in call to
    property_entry_read_int_array() to retrieve the correct nargs value.

    Fixes: b06184acf751 ("software node: Add software_node_get_reference_args()")
    Signed-off-by: Clément Léger <clement.leger@bootlin.com>
    Reviewed-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Reviewed-by: Daniel Scally <djrscally@gmail.com>
    Acked-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41194f2f36dddfc61214bcbf35e8686b42a5cd1d
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Mon Nov 15 21:34:59 2021 +0100

    backlight: qcom-wled: Respect enabled-strings in set_brightness

    [ Upstream commit ec961cf3241153e0f27d850f1bf0f172e7d27a21 ]

    The hardware is capable of controlling any non-contiguous sequence of
    LEDs specified in the DT using qcom,enabled-strings as u32
    array, and this also follows from the DT-bindings documentation.  The
    numbers specified in this array represent indices of the LED strings
    that are to be enabled and disabled.

    Its value is appropriately used to setup and enable string modules, but
    completely disregarded in the set_brightness paths which only iterate
```

```
    over the number of strings linearly.
    Take an example where only string 2 is enabled with
    qcom,enabled_strings=<2>: this string is appropriately enabled but
    subsequent brightness changes would have only touched the zero'th
    brightness register because num_strings is 1 here.  This is simply
    addressed by looking up the string for this index in the enabled_strings
    array just like the other codepaths that iterate over num_strings.

    Likewise enabled_strings is now also used in the autodetection path for
    consistent behaviour: when a list of strings is specified in DT only
    those strings will be probed for autodetection, analogous to how the
    number of strings that need to be probed is already bound by
    qcom,num-strings.  After all autodetection uses the set_brightness
    helpers to set an initial value, which could otherwise end up changing
    brightness on a different set of strings.

    Fixes: 775d2ffb4af6 ("backlight: qcom-wled: Restructure the driver for WLED3")
    Fixes: 03b2b5e86986 ("backlight: qcom-wled: Add support for WLED4 peripheral")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-10-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cee8d5c45c2e5b83630c8947196b619ac6996e6
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Mon Nov 15 21:34:53 2021 +0100

    backlight: qcom-wled: Use cpu_to_le16 macro to perform conversion

    [ Upstream commit 0a139358548968b2ff308257b4fbeec7badcc3e1 ]

    The kernel already provides appropriate primitives to perform endianness
    conversion which should be used in favour of manual bit-wrangling.

    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-4-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 585bf131dff85fcb8f08f9f8e2168115ab77bc25
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Mon Nov 15 21:34:55 2021 +0100

    backlight: qcom-wled: Override default length with qcom,enabled-strings

    [ Upstream commit 2b4b49602f9feca7b7a84eaa33ad9e666c8aa695 ]

    The length of qcom,enabled-strings as property array is enough to
    determine the number of strings to be enabled, without needing to set
    qcom,num-strings to override the default number of strings when less
    than the default (which is also the maximum) is provided in DT.

    This also introduces an extra warning when qcom,num-strings is set,
    denoting that it is not necessary to set both anymore.  It is usually
    more concise to set just qcom,num-length when a zero-based, contiguous
    range of strings is needed (the majority of the cases), or to only set
    qcom,enabled-strings when a specific set of indices is desired.

    Fixes: 775d2ffb4af6 ("backlight: qcom-wled: Restructure the driver for WLED3")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
```

```
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-6-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 007cfb6373595ab3a52d3f8a7a492fc3541812e2
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Mon Nov 15 21:34:54 2021 +0100

    backlight: qcom-wled: Fix off-by-one maximum with default num_strings

    [ Upstream commit 5ada78b26f935f8751852dffa24f6b545b1d2517 ]

    When not specifying num-strings in the DT the default is used, but +1 is
    added to it which turns WLED3 into 4 and WLED4/5 into 5 strings instead
    of 3 and 4 respectively, causing out-of-bounds reads and register
    read/writes.  This +1 exists for a deficiency in the DT parsing code,
    and is simply omitted entirely - solving this oob issue - by parsing the
    property separately much like qcom,enabled-strings.

    This also enables more stringent checks on the maximum value when
    qcom,enabled-strings is provided in the DT, by parsing num-strings after
    enabled-strings to allow it to check against (and in a subsequent patch
    override) the length of enabled-strings: it is invalid to set
    num-strings higher than that.
    The DT currently utilizes it to get around an incorrect fixed read of
    four elements from that array (has been addressed in a prior patch) by
    setting a lower num-strings where desired.

    Fixes: 93c64f1ea1e8 ("leds: add Qualcomm PM8941 WLED driver")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-By: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-5-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4397f8dcd04c4be5e992c7325842027a7a19fc3
Author: Marijn Suijten <marijn.suijten@somainline.org>
Date:   Mon Nov 15 21:34:52 2021 +0100

    backlight: qcom-wled: Pass number of elements to read to read_u32_array

    [ Upstream commit e29e24bdabfeddbf8b1a4ecac1af439a85150438 ]

    of_property_read_u32_array takes the number of elements to read as last
    argument. This does not always need to be 4 (sizeof(u32)) but should
    instead be the size of the array in DT as read just above with
    of_property_count_elems_of_size.

    To not make such an error go unnoticed again the driver now bails
    accordingly when of_property_read_u32_array returns an error.
    Surprisingly the indentation of newlined arguments is lining up again
    after prepending `rc = `.

    Fixes: 775d2ffb4af6 ("backlight: qcom-wled: Restructure the driver for WLED3")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-3-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f04c0d3288991fd950a4733ba28b4bbc4ea38b3
Author: Marijn Suijten <marijn.suijten@somainline.org>
```

```
Date:   Mon Nov 15 21:34:51 2021 +0100

    backlight: qcom-wled: Validate enabled string indices in DT

    [ Upstream commit c05b21ebc5bce3ecc78c2c71afd76d92c790a2ac ]

    The strings passed in DT may possibly cause out-of-bounds register
    accesses and should be validated before use.

    Fixes: 775d2ffb4af6 ("backlight: qcom-wled: Restructure the driver for WLED3")
    Signed-off-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@somainline.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211115203459.1634079-2-marijn.suijten@somainline.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92f5df7f52405a6acb639a4ab12f7ac73d5f0de8
Author: Paul Chaignon <paul@isovalent.com>
Date:   Mon Dec 20 22:45:28 2021 +0100

    bpftool: Enable line buffering for stdout

    [ Upstream commit 1a1a0b0364ad291bd8e509da104ac8b5b1afec5d ]

    The output of bpftool prog tracelog is currently buffered, which is
    inconvenient when piping the output into other commands. A simple
    tracelog | grep will typically not display anything. This patch fixes it
    by enabling line buffering on stdout for the whole bpftool binary.

    Fixes: 30da46b5dc3a ("tools: bpftool: add a command to dump the trace pipe")
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Paul Chaignon <paul@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Link: https://lore.kernel.org/bpf/20211220214528.GA11706@Mem
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a44c79dba47acb9ae2165557e2681253833eb5f8
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Dec 8 15:35:48 2021 -0800

    Bluetooth: L2CAP: Fix using wrong mode

    [ Upstream commit 30d57722732d9736554f85f75f9d7ad5402d192e ]

    If user has a set to use SOCK_STREAM the socket would default to
    L2CAP_MODE_ERTM which later needs to be adjusted if the destination
    address is LE which doesn't support such mode.

    Fixes: 15f02b9105625 ("Bluetooth: L2CAP: Add initial code for Enhanced Credit Based
Mode")
    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 99ac307d9826d60d7b9ec91487acf5981fe3a27a
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Thu Sep 16 13:09:22 2021 +0200

    um: virtio_uml: Fix time-travel external time propagation

    [ Upstream commit 85e73968a040c642fd38f6cba5b73b61f5d0f052 ]

    When creating an external event, the current time needs to
```

```
    be propagated to other participants of a simulation. This
    is done in the places here where we kick a virtq etc.

    However, it must be done for _all_ external events, and
    that includes making the initial socket connection and
    later closing it. Call time_travel_propagate_time() to do
    this before making or closing the socket connection.

    Apparently, at least for the initial connection creation,
    due to the remote side in my use cases using microseconds
    (rather than nanoseconds), this wasn't a problem yet; only
    started failing between 5.14-rc1 and 5.15-rc1 (didn't test
    others much), or possibly depending on the configuration,
    where more delays happen before the virtio devices are
    initialized.

    Fixes: 88ce64249233 ("um: Implement time-travel=ext")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c2ecc54231352f2ed4ef8f37bd4c4a59b78c26e
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Sep 15 20:30:22 2021 +0200

    lib/logic_iomem: Fix operation on 32-bit

    [ Upstream commit 4e8a5edac5010820e7c5303fc96f5a262e096bb6 ]

    On 32-bit, the first entry might be at 0/NULL, but that's
    strange and leads to issues, e.g. where we check "if (ret)".
    Use a IOREMAP_BIAS/IOREMAP_MASK of 0x80000000UL to avoid
    this. This then requires reducing the number of areas (via
    MAX_AREAS), but we still have 128 areas, which is enough.

    Fixes: ca2e334232b6 ("lib: add iomem emulation (logic_iomem)")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cbf77de7afce88a067306abfee370d913b57f01
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Sep 15 20:30:21 2021 +0200

    lib/logic_iomem: Fix 32-bit build

    [ Upstream commit 4e84139e14af5ea60772cc4f33d7059aec76e0eb ]

    On a 32-bit build, the (unsigned long long) casts throw warnings
    (or errors) due to being to a different integer size. Cast to
    uintptr_t first (with the __force for sparse) and then further
    to get the consistent print on 32 and 64-bit.

    Fixes: ca2e334232b6 ("lib: add iomem emulation (logic_iomem)")
    Reported-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit af2668c3d912fd4c638b49d212969a60c6af9894
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Sep 15 20:30:20 2021 +0200

    um: virt-pci: Fix 32-bit compile
```

```
    [ Upstream commit d73820df6437b5d0a57be53faf39db46a0264b3a ]

    There were a few 32-bit compile warnings that of course
    turned into errors with -Werror, fix the 32-bit build.

    Fixes: 68f5d3f3b654 ("um: add PCI over virtio emulation driver")
    Reported-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c0d5404dd2b43c2dfa756ffdf78c136f028197b
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Tue Aug 31 09:11:15 2021 +0200

    um: rename set_signals() to um_set_signals()

    [ Upstream commit bbe33504d4a7fdab9011211e55e262c869b3f6cc ]

    Rename set_signals() as there's at least one driver that
    uses the same name and can now be built on UM due to PCI
    support, and thus we can get symbol conflicts.

    Also rename set_signals_trace() to be consistent.

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: 68f5d3f3b654 ("um: add PCI over virtio emulation driver")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f3df33728dcdd26f56a6db00b11025852949fc6
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Tue Aug 31 09:27:53 2021 +0200

    um: fix ndelay/udelay defines

    [ Upstream commit 5f8539e2ff962e25b57742ca7106456403abbc94 ]

    Many places in the kernel use 'udelay' as an identifier, and
    are broken with the current "#define udelay um_udelay". Fix
    this by adding an argument to the macro, and do the same to
    'ndelay' as well, just in case.

    Fixes: 0bc8fb4dda2b ("um: Implement ndelay/udelay in time-travel mode")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e09744d90bfdf37a6215a7b1f56ceda6b38ed50
Author: Bernard Zhao <bernard@vivo.com>
Date:   Fri Dec 10 04:03:58 2021 -0800

    selinux: fix potential memleak in selinux_add_opt()

    [ Upstream commit 2e08df3c7c4e4e74e3dd5104c100f0bf6288aaa8 ]

    This patch try to fix potential memleak in error branch.

    Fixes: ba6418623385 ("selinux: new helper - selinux_add_opt()")
    Signed-off-by: Bernard Zhao <bernard@vivo.com>
    [PM: tweak the subject line, add Fixes tag]
    Signed-off-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit c4e1577ccb7b8f6917144a63931960923a236453
Author: Christoph Hellwig <hch@lst.de>
Date:   Tue Dec 21 17:18:51 2021 +0100

    block: fix error unwinding in device_add_disk

    [ Upstream commit 99d8690aae4b2f0d1d90075de355ac087f820a66 ]

    One device_add is called disk->ev will be freed by disk_release, so we
    should free it twice.  Fix this by allocating disk->ev after device_add
    so that the extra local unwinding can be removed entirely.

    Based on an earlier patch from Tetsuo Handa.

    Reported-by: syzbot <syzbot+28a66a9fbc621c939000@syzkaller.appspotmail.com>
    Tested-by: syzbot <syzbot+28a66a9fbc621c939000@syzkaller.appspotmail.com>
    Fixes: 83cbce9574462c6b ("block: add error handling for device_add_disk / add_disk")
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20211221161851.788424-1-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 454c33da03ce69b588bc074c1a1e9f9b8e4302d5
Author: Sergey Shtylyov <s.shtylyov@omp.ru>
Date:   Fri Dec 17 23:27:17 2021 +0300

    mmc: meson-mx-sdio: add IRQ check

    [ Upstream commit 8fc9a77bc64e1f23d07953439817d8402ac9706f ]

    The driver neglects to check the result of platform_get_irq()'s call and
    blithely passes the negative error codes to devm_request_threaded_irq()
    (which takes *unsigned* IRQ #), causing it to fail with -EINVAL, overriding
    an original error code. Stop calling devm_request_threaded_irq() with the
    invalid IRQ #s.

    Fixes: ed80a13bb4c4 ("mmc: meson-mx-sdio: Add a driver for the Amlogic Meson8 and
Meson8b SoC")
    Signed-off-by: Sergey Shtylyov <s.shtylyov@omp.ru>
    Reviewed-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Link: https://lore.kernel.org/r/20211217202717.10041-3-s.shtylyov@omp.ru
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 60d62caf6e1d6a5f4419ef901a987c6321d4b3fe
Author: Sergey Shtylyov <s.shtylyov@omp.ru>
Date:   Fri Dec 17 23:27:16 2021 +0300

    mmc: meson-mx-sdhc: add IRQ check

    [ Upstream commit 77bed755e0f06135faccdd3948863703f9a6e640 ]

    The driver neglects to check the result of platform_get_irq()'s call and
    blithely passes the negative error codes to devm_request_threaded_irq()
    (which takes *unsigned* IRQ #), causing it to fail with -EINVAL, overriding
    an original error code. Stop calling devm_request_threaded_irq() with the
    invalid IRQ #s.

    Fixes: e4bf1b0970ef ("mmc: host: meson-mx-sdhc: new driver for the Amlogic Meson SDHC
host")
    Signed-off-by: Sergey Shtylyov <s.shtylyov@omp.ru>
    Reviewed-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Link: https://lore.kernel.org/r/20211217202717.10041-2-s.shtylyov@omp.ru
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
```

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

commit 0696773fcc2806fa737d17015b15154f571067f6
Author: Avraham Stern <avraham.stern@intel.com>
Date:   Sun Dec 19 13:28:26 2021 +0200

    iwlwifi: mvm: set protected flag only for NDP ranging

    [ Upstream commit 6bb2ea37c02db98cb677f978cfcb833ca608c5eb ]

    Don't use protected ranging negotiation for FTM ranging as responders
    that support only FTM ranging don't expect the FTM request to be
    protected.

    Signed-off-by: Avraham Stern <avraham.stern@intel.com>
    Fixes: 517a5eb9fab2 ("iwlwifi: mvm: when associated with PMF, use protected NDP ranging
negotiation")
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219132536.f50ed0e3c6b3.Ibff247ee9d4e6e0a1a2d08a3c8a4b
bb37e6829dd@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f24be5e7e86d76981900663e2d7898480574b15
Author: Avraham Stern <avraham.stern@intel.com>
Date:   Sun Dec 19 12:18:20 2021 +0200

    iwlwifi: mvm: perform 6GHz passive scan after suspend

    [ Upstream commit f4745cbb17572209a7fa27a6796ed70e7ada860b ]

    The 6GHz passive scan is performed only once every 50 minutes.
    However, in case of suspend/resume, the regulatory information
    is reset, so 6GHz channels may become disabled.
    Fix it by performing a 6GHz passive scan within 60 seconds after
    suspend/resume even if the 50 minutes timeout did not expire yet.

    Signed-off-by: Avraham Stern <avraham.stern@intel.com>
    Fixes: e8fe3b41c3a3 ("iwlwifi: mvm: Add support for 6GHz passive scan")
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219121514.6d5c043372cf.I251dd5618a3f0b8febbcca788eb86
1f1cd6039bc@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b00f80dbc5eb15b9477ecf0dd5f996433fc0186
Author: Nathan Errera <nathan.errera@intel.com>
Date:   Sun Dec 19 12:18:15 2021 +0200

    iwlwifi: mvm: test roc running status bits before removing the sta

    [ Upstream commit 998e1aba6e5eb35370eaf30ccc1823426ec11f90 ]

    In some cases the sta is being removed twice since we do not test the
    roc aux running before removing it. Start looking at the bit before
    removing the sta.

    Signed-off-by: Nathan Errera <nathan.errera@intel.com>
    Fixes: 2c2c3647cde4 ("iwlwifi: mvm: support ADD_STA_CMD_API_S ver 12")
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219121514.d5376ac6bcb0.Ic5f8470ea60c072bde9d1503e5f52
8b65e301e20@changeid
```

```
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1eb83e559b4b10d7474d3bd81fc9aba3fd98ed74
Author: Luca Coelho <luciano.coelho@intel.com>
Date:   Sun Dec 19 12:18:14 2021 +0200

    iwlwifi: don't pass actual WGDS revision number in table_revision

    [ Upstream commit ac9952f6954216be72a8bde210e1ef2c949d8ee0 ]

    The FW API for PER_CHAIN_LIMIT_OFFSET_CMD is misleading.  The element
    name is table_rev, but it shouldn't actually contain the table
    revision number, but whether we should use the South Korea scheme or
    not.

    Fix the driver so that we only set this value to either 0 or 1.  It
    will only be 1 (meaning South Korea) if the ACPI WGDS table revision
    is 1.

    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Fixes: 664c011b763e ("iwlwifi: acpi: support reading and storing WGDS revision 2")
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Link:
https://lore.kernel.org/r/iwlwifi.20211219121514.abed3b8119c7.I1fdc2c14577523fcffdfe8fb5902c
2d8efde7e09@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 381a0452dd311f3a7d5ee1e780ffd4b14702a2c5
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Sun Dec 19 11:14:18 2021 +0200

    iwlwifi: mvm: fix 32-bit build in FTM

    [ Upstream commit 8b0f92549f2c2458200935c12a2e2a6e80234cf5 ]

    On a 32-bit build, the division here needs to be done
    using do_div(), otherwise the compiler will try to call
    a function that doesn't exist, thus failing to build.

    Fixes: b68bd2e3143a ("iwlwifi: mvm: Add FTM initiator RTT smoothing logic")
    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.20211219111352.e56cbf614a4d.Ib98004ccd2c7a55fd883a8ea7eebd
810f406dec6@changeid
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f21e8cb927fa9da49490b74ddda3856a828b82d6
Author: Kai-Heng Feng <kai.heng.feng@canonical.com>
Date:   Wed Dec 15 19:46:34 2021 +0800

    rtw88: Disable PCIe ASPM while doing NAPI poll on 8821CE

    [ Upstream commit 24f5e38a13b5ae2b6105cda8bb47c19108e62a9a ]

    Many Intel based platforms face system random freeze after commit
    9e2fd29864c5 ("rtw88: add napi support").

    The commit itself shouldn't be the culprit. My guess is that the 8821CE
    only leaves ASPM L1 for a short period when IRQ is raised. Since IRQ is
    masked during NAPI polling, the PCIe link stays at L1 and makes RX DMA
    extremely slow. Eventually the RX ring becomes messed up:
```

```
      [ 1133.194697] rtw_8821ce 0000:02:00.0: pci bus timeout, check dma status

      Since the 8821CE hardware may fail to leave ASPM L1, manually do it in
      the driver to resolve the issue.

      Fixes: 9e2fd29864c5 ("rtw88: add napi support")
      Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=215131
      BugLink: https://bugs.launchpad.net/bugs/1927808
      Signed-off-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
      Acked-by: Jian-Hong Pan <jhp@endlessos.org>
      Signed-off-by: Kalle Valo <kvalo@kernel.org>
      Link: https://lore.kernel.org/r/20211215114635.333767-1-kai.heng.feng@canonical.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit db7e4381c811fe372946cd5636cbca66ccfde9d3
  Author: Ping-Ke Shih <pkshih@realtek.com>
  Date:   Fri Nov 19 13:24:37 2021 +0800

      rtw88: add quirk to disable pci caps on HP 250 G7 Notebook PC

      [ Upstream commit c81edb8dddaa36c4defa26240cc19127f147283f ]

      8821CE causes random freezes on HP 250 G7 Notebook PC. Add a quirk
      to disable pci ASPM capability.

      Reported-by: rtl8821cerfe2 <rtl8821cerfe2@protonmail.com>
      Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
      Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
      Link: https://lore.kernel.org/r/20211119052437.8671-1-pkshih@realtek.com
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 61371a716228f8788ca6c40cc2668aaacac9b9f6
  Author: Dan Carpenter <dan.carpenter@oracle.com>
  Date:   Fri Dec 17 18:03:12 2021 +0300

      wilc1000: fix double free error in probe()

      [ Upstream commit 4894edacfa93d7046bec4fc61fc402ac6a2ac9e8 ]

      Smatch complains that there is a double free in probe:

      drivers/net/wireless/microchip/wilc1000/spi.c:186 wilc_bus_probe() error: double free of
  'spi_priv'
      drivers/net/wireless/microchip/wilc1000/sdio.c:163 wilc_sdio_probe() error: double free
  of 'sdio_priv'

      The problem is that wilc_netdev_cleanup() function frees "wilc->bus_data".
      That's confusing and a layering violation.  Leave the frees in probe(),
      delete the free in wilc_netdev_cleanup(), and add some new frees to the
      remove() functions.

      Fixes: dc8b338f3bcd ("wilc1000: use goto labels on error path")
      Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
      Reviewed-by: Claudiu Beznea <claudiu.beznea@microchip.com>
      Signed-off-by: Kalle Valo <kvalo@kernel.org>
      Link: https://lore.kernel.org/r/20211217150311.GC16611@kili
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 035bf9c53ef1c9248d2eba3a859809b6932a8d7d
  Author: Sean Wang <sean.wang@mediatek.com>
  Date:   Thu Dec 16 05:25:37 2021 +0800

      mt76: mt7921s: fix suspend error with enlarging mcu timeout value

      [ Upstream commit 1bb42a354d8ca2888c7c2fcbf0add410176a33dc ]
```

```
    Fix the false positive suspend error that may occur on mt7921s
    with enlarging mcu timeout value.

    The reason why we have to enlarge mcu timeout from HZ / 3 to HZ is
    we should consider the additional overhead caused by running
    concurrently with btmtksdio (a MT7921 bluetooth SDIO driver) that
    would compete for the same SDIO bus in process context to complete
    the suspend procedure.

    Fixes: 48fab5bbef40 ("mt76: mt7921: introduce mt7921s support")
    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 a38b94c43943286a1564cd792efc668da5b262c1
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Dec 9 14:06:26 2021 +0100

    mt76: connac: introduce MCU_UNI_CMD macro

    [ Upstream commit 5472240245793c13e9986c61dd34c697296deed4 ]

    Similar to MCU_EXT_CMD, introduce MCU_UNI_CMD for unified commands

    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 b8598139a827d370dfcd5c2539f9f224cd138b14
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Dec 9 14:06:25 2021 +0100

    mt76: connac: remove MCU_FW_PREFIX bit

    [ Upstream commit 7159eb828d21ad8fb3c1c1782a286a658ba6bf66 ]

    Get rid of MCU_FW_PREFIX bit since it is no longer used

    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 8b103fae4f075c854ecea6425641df6d7a45fc92
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Dec 9 14:06:24 2021 +0100

    mt76: connac: align MCU_EXT definitions with 7915 driver

    [ Upstream commit 9d8d136cf0b6d5578442f38ea9eefdf67cc84fc4 ]

    Align MCU_EXT and MCU_FW definitions between mt76_connac and mt7915
    driver. This is a preliminary patch to reuse mt76_connac in mt7915
    driver.

    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 ec31853741f6a3ba281f1afbf1f2d087c24afc60
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Dec 9 14:06:23 2021 +0100

    mt76: connac: introduce MCU_EXT macros
```

```
    [ Upstream commit e6d2070d9d64aad04c12424865cfd9684ba64bea ]

    Introduce MCU_EXT_CMD and MCU_EXT_QUERY macros in mt76_connac module.
    This is a preliminary patch to reuse mt76_connac module in mt7915
    driver.

    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 c78938707accaeb19383d9a0f92ea4807b10360e
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Thu Dec 16 05:25:35 2021 +0800

    mt76: mt7921: fix possible resume failure

    [ Upstream commit 5375001bb4ce22801bf3bb566cc3e67d2d3a5dc0 ]

    Fix the possible resume failure due to mt76_connac_mcu_set_hif_suspend
    timeout.

    That is because clearing the flag pm->suspended too early opened up a race
    window, where mt7921_poll_tx/rx scheduled a ps_work to put the device in
    doze mode, that is unexpected for the device is being resumed from the
    suspend state and would make the remaining MCU comamnds in resume handler
    failed to execute.

    Fixes: ffa1bf97425b ("mt76: mt7921: introduce PM support")
    Co-developed-by: YN Chen <YN.Chen@mediatek.com>
    Signed-off-by: YN Chen <YN.Chen@mediatek.com>
    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 cec4a2388a1ab8911a1b28b4d88faab4e34dd4ab
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Tue Nov 30 10:36:15 2021 +0100

    mt76: connac: fix last_chan configuration in mt76_connac_mcu_rate_txpower_band

    [ Upstream commit 73c7c0443685d8d152c3451e7305a2c2bc47b32e ]

    last_ch configuration must not be dependent on the current configured band
    but it is defined by hw capabilities since the fw always expects the
    following order:
    - 2GHz
    - 5GHz
    - 6GHz

    Fixes: 9b2ea8eee42a1 ("mt76: connac: set 6G phymode in single-sku support")
    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 1542711d31f014597e5f343236614ed1621b52b3
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Tue Nov 30 07:05:28 2021 +0800

    mt76: mt7921s: fix possible kernel crash due to invalid Rx count

    [ Upstream commit 2b7f3574ca9a7ff4a6b4ec1ae4dfdfde481ac80b ]

    Return the proper error code when out-of-range the Rx aggregation count
    are reported from the hardware that would create the unreasonable extreme
    large Rx buffer.
```

```
    [  100.873810]  show_stack+0x20/0x2c
    [  100.873823]  dump_stack+0xc4/0x140
    [  100.873839]  bad_page+0x110/0x114
    [  100.873854]  check_new_pages+0xf8/0xfc
    [  100.873869]  rmqueue+0x5a0/0x640
    [  100.873884]  get_page_from_freelist+0x124/0x20c
    [  100.873898]  __alloc_pages_nodemask+0x114/0x2a4
    [  100.873918]  mt76s_rx_run_queue+0xd4/0x2e4 [mt76_sdio
8280a88a0c8c9cf203f16e194f99ac293bdbb2f5]
    [  100.873938]  mt76s_rx_handler+0xd4/0x2a0 [mt76_sdio
8280a88a0c8c9cf203f16e194f99ac293bdbb2f5]
    [  100.873957]  mt76s_txrx_worker+0xac/0x17c [mt76_sdio
8280a88a0c8c9cf203f16e194f99ac293bdbb2f5]
    [  100.873977]  mt7921s_txrx_worker+0x5c/0xd8 [mt7921s
d0bdbc018082dbc8dc1407614be3c2e7bd64423b]
    [  100.874003]  __mt76_worker_fn+0xe8/0x170 [mt76
b80af3483a8f9d48e916c12d8dbfaa0d3cd15337]
    [  100.874018]  kthread+0x148/0x3ac
    [  100.874032]  ret_from_fork+0x10/0x30
    [  100.874067] Kernel Offset: 0x1fe2000000 from 0xffffffc010000000
    [  100.874079] PHYS_OFFSET: 0xffffffe800000000
    [  100.874090] CPU features: 0x0240002,2188200c

    Fixes: 48fab5bbef40 ("mt76: mt7921: introduce mt7921s support")
    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 1572ad73f11eaeb6dc26f8332c7bdf9d3ffeac2f
Author: Shayne Chen <shayne.chen@mediatek.com>
Date:   Sat Nov 27 00:53:16 2021 +0800

    mt76: mt7921: use correct iftype data on 6GHz cap init

    [ Upstream commit 00ff52346d74c38787ff8b4acde8c5671d9b7fe2 ]

    Set 6GHz cap to iftype data which is matched to the type of
    current interface.

    Fixes: 50ac15a511e3 ("mt76: mt7921: add 6GHz support")
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Acked-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6bd23dc9eabf0ff790ff20c3ce8ce627c9f676ac
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Sat Nov 20 07:22:11 2021 +0800

    mt76: mt7921s: fix the device cannot sleep deeply in suspend

    [ Upstream commit 5ad4faca7690a88b4529238c757b6b8ead8056ec ]

    According to the MT7921S firmware, the cmd MCU_UNI_CMD_HIF_CTRL have to
    be last MCU command to execute in suspend handler and all data traffic
    have to be stopped before the cmd MCU_UNI_CMD_HIF_CTRL starts as well
    in order that mt7921 can successfully fall into the deep sleep mode.

    Where we reuse the flag MT76_STATE_SUSPEND and avoid creating
    another global flag to stop all of the traffic onto the SDIO bus.

    Fixes: 48fab5bbef40 ("mt76: mt7921: introduce mt7921s support")
    Reported-by: Leon Yen <leon.yen@mediatek.com>
    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 bae9b748c7713ec56b4d0e020f4909595dc9b94e
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Sat Nov 20 07:22:10 2021 +0800

    mt76: mt7921: move mt76_connac_mcu_set_hif_suspend to bus-related files

    [ Upstream commit 6906aa93eb93d54a42ce1902f00d6ea04ecb039b ]

    This is a preliminary patch for the following patch
    ("mt76: mt7921s: fix the device cannot sleep deeply in suspend).

    mt76_connac_mcu_set_hif_suspend eventually would be handled in each
    bus-level suspend/resume handler in either mt7921/sdio.c or mt7921/pci.c
    depending on what type of the bus the device is running on. We can move
    mt76_connac_mcu_set_hif_suspend to bus-related files to simplify the logic.

    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 633666158f46d958aaed60a4ec641c599d89275e
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Nov 18 11:50:27 2021 +0100

    mt76: mt7921: fix possible NULL pointer dereference in mt7921_mac_write_txwi

    [ Upstream commit ec2ebc1c5a5ce49538fa78108c53eaf722eee908 ]

    Fix a possible NULL pointer deference issue in mt7921_mac_write_txwi
    routine if vif is NULL.

    Fixes: 33920b2bf0483 ("mt76: add support for setting mcast rate")
    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 04d18e8af7a5ed4da5d49dfd11dd6544db6466ce
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Thu Nov 18 11:42:06 2021 +0100

    mt76: fix possible OOB issue in mt76_calculate_default_rate

    [ Upstream commit d4f3d1c4d3c2bcce76a96a6562170664b25112f0 ]

    Cap max offset value to ARRAY_SIZE(mt76_rates) - 1 in
    mt76_calculate_default_rate routine in order to avoid possible Out Of
    Bound accesses.

    Fixes: 33920b2bf0483 ("mt76: add support for setting mcast rate")
    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 ea45a71ef323602d769f62fb907f4ff29d870b9a
Author: Lorenzo Bianconi <lorenzo@kernel.org>
Date:   Mon Nov 8 11:03:52 2021 +0100

    mt76: debugfs: fix queue reporting for mt76-usb

    [ Upstream commit eae7df016c30468e09481a75b1339f0ab5ece222 ]

    Fix number of rx-queued frames reported by mt76_usb driver.
```

```
    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: 2d8be76c1674 ("mt76: debugfs: improve queue node readability")
    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 cdb39e251f864910b2fb6c099b1ef9d12c6e22c7
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Wed Nov 3 08:05:16 2021 +0800

    mt76: mt7921: fix MT7921E reset failure

    [ Upstream commit 0efaf31dec572d3aac4316c6d952e06d1c33adc4 ]

    There is a missing mt7921e_driver_own in the MT7921E reset procedure
    since the mt7921 mcu.c has been refactored for MT7921S, that will
    result in MT7921E reset failure, so add it back now.

    Fixes: dfc7743de1eb ("mt76: mt7921: refactor mcu.c to be bus independent")
    Reported-by: YN Chen <YN.Chen@mediatek.com>
    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 50e7e4db80a1388f0b89df112ba6e1303c09cc7e
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Fri Oct 29 06:29:58 2021 +0800

    mt76: mt7921: drop offload_flags overwritten

    [ Upstream commit 2363b6a646b65a207345b9a9024dff0eff3fec44 ]

    offload_flags have to be dropped for mt7921 because mt76.omac_idx 0 would
    always run as station mode that would cause Tx encapsulation setting
    is not applied to mac80211.

    Also, drop IEEE80211_OFFLOAD_ENCAP_4ADDR too because it is not really
    being supported.

    Fixes: e0f9fdda81bd ("mt76: mt7921: add ieee80211_ops")
    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 832c2a994a0344190fd5804d3ab0c5a63636691f
Author: Marek Behún <kabel@kernel.org>
Date:   Tue Nov 9 17:46:04 2021 +0100

    ARM: dts: armada-38x: Add generic compatible to UART nodes

    [ Upstream commit 62480772263ab6b52e758f2346c70a526abd1d28 ]

    Add generic compatible string "ns16550a" to serial port nodes of Armada
    38x.

    This makes it possible to use earlycon.

    Fixes: 0d3d96ab0059 ("ARM: mvebu: add Device Tree description of the Armada 380/385
SoCs")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit b21b215d3205bb0d86dc76c6540289d41079dddc
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Fri Nov 12 14:44:03 2021 +0100

    arm64: dts: marvell: cn9130: enable CP0 GPIO controllers

    [ Upstream commit 0734f8311ce72c9041e5142769eff2083889c172 ]

    CN9130 has a built-in CP115 which has 2 GPIO controllers, but unlike in
    Armada 7k and 8k both are left disabled by the SoC DTSI.

    This first of all makes no sense as they are always present due to being
    SoC built-in and its an issue as boards like CN9130-CRB use the CP0 GPIO2
    pins for regulators and SD card support without enabling them first.

    So, enable both of them like Armada 7k and 8k do.

    Fixes: 6b8970bd8d7a ("arm64: dts: marvell: Add support for Marvell CN9130 SoC support")

    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a158961c45603a18f371356a94206f4e2004102
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Fri Nov 12 14:44:02 2021 +0100

    arm64: dts: marvell: cn9130: add GPIO and SPI aliases

    [ Upstream commit effd42600b987c1e95f946b14fefc1c7639e7439 ]

    CN9130 has one CP115 built in, which like the CP110 has 2 GPIO and 2 SPI
    controllers built-in.

    However, unlike the Armada 7k and 8k the SoC DTSI doesn't add the required
    aliases as both the Orion SPI driver and MVEBU GPIO drivers require the
    aliases to be present.

    So add the required aliases for GPIO and SPI controllers.

    Fixes: 6b8970bd8d7a ("arm64: dts: marvell: Add support for Marvell CN9130 SoC support")

    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32ad93587357a0a88338b21f6e3931216f3d578c
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Fri Dec 17 16:34:28 2021 +0800

    usb: ftdi-elan: fix memory leak on device disconnect

    [ Upstream commit 1646566b5e0c556f779180a8514e521ac735de1e ]

    'ftdi' is alloced when probe device, but not free on device disconnect,
    this cause a memory leak as follows:

    unreferenced object 0xffff88800d584000 (size 8400):
      comm "kworker/0:2", pid 3809, jiffies 4295453055 (age 13.784s)
      hex dump (first 32 bytes):
        00 40 58 0d 80 88 ff ff 00 40 58 0d 80 88 ff ff  .@X......@X.....
        00 00 00 00 00 00 00 00 00 00 00 ad 4e ad de  ............N..
      backtrace:
        [<000000000d47f947>] kmalloc_order_trace+0x19/0x110 mm/slab_common.c:960
        [<000000008548ac68>] ftdi_elan_probe+0x8c/0x880 drivers/usb/misc/ftdi-elan.c:2647
```

```
         [<000000007f73e422>] usb_probe_interface+0x31b/0x800 drivers/usb/core/driver.c:396
         [<00000000fe8d07fc>] really_probe+0x299/0xc30 drivers/base/dd.c:517
         [<0000000005da7d32>] __driver_probe_device+0x357/0x500 drivers/base/dd.c:751
         [<000000003c2c9579>] driver_probe_device+0x4e/0x140 drivers/base/dd.c:781

    Fix it by freeing 'ftdi' after nobody use it.

    Fixes: a5c66e4b2418 ("USB: ftdi-elan: client driver for ELAN Uxxx adapters")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Link: https://lore.kernel.org/r/20211217083428.2441-1-weiyongjun1@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 01669a79465fb99a14abbbc84bdb03e40aaf54b1
Author: Andre Przywara <andre.przywara@arm.com>
Date:   Mon Nov 22 16:28:43 2021 +0100

    ARM: 9159/1: decompressor: Avoid UNPREDICTABLE NOP encoding

    [ Upstream commit a92882a4d270fbcc021ee6848de5e48b7f0d27f3 ]

    In the decompressor's head.S we need to start with an instruction that
    is some kind of NOP, but also mimics as the PE/COFF header, when the
    kernel is linked as an UEFI application. The clever solution here is
    "tstne r0, #0x4d000", which in the worst case just clobbers the
    condition flags, and bears the magic "MZ" signature in the lowest 16 bits.

    However the encoding used (0x13105a4d) is actually not valid, since bits
    [15:12] are supposed to be 0 (written as "(0)" in the ARM ARM).
    Violating this is UNPREDICTABLE, and *can* trigger an UNDEFINED
    exception. Common Cortex cores seem to ignore those bits, but QEMU
    chooses to trap, so the code goes fishing because of a missing exception
    handler at this point. We are just saved by the fact that commonly (with
    -kernel or when running from U-Boot) the "Z" bit is set, so the
    instruction is never executed. See [0] for more details.

    To make things more robust and avoid UNPREDICTABLE behaviour in the
    kernel code, lets replace this with a "two-instruction NOP":
    The first instruction is an exclusive OR, the effect of which the second
    instruction reverts. This does not leave any trace, neither in a
    register nor in the condition flags. Also it's a perfectly valid
    encoding. Kudos to Peter Maydell for coming up with this gem.

    [0] https://lore.kernel.org/qemu-devel/YTPIdbUCmwagL5%2FD@os.inf.tu-dresden.de/T/

    Link: https://lore.kernel.org/linux-arm-kernel/20210908162617.104962-1-
andre.przywara@arm.com/T/

    Fixes: 81a0bc39ea19 ("ARM: add UEFI stub support")
    Signed-off-by: Andre Przywara <andre.przywara@arm.com>
    Reported-by: Adam Lackorzynski <adam@l4re.org>
    Suggested-by: Peter Maydell <peter.maydell@linaro.org>
    Reviewed-by: Ard Biesheuvel <ardb@kernel.org>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7138341e713220360936d4d997865de431bb2a2b
Author: Antony Antony <antony.antony@secunet.com>
Date:   Sun Dec 12 11:35:00 2021 +0100

    xfrm: state and policy should fail if XFRMA_IF_ID 0

    [ Upstream commit 68ac0f3810e76a853b5f7b90601a05c3048b8b54 ]
```

```
    xfrm ineterface does not allow xfrm if_id = 0
    fail to create or update xfrm state and policy.

    With this commit:
     ip xfrm policy add src 192.0.2.1 dst 192.0.2.2 dir out if_id 0
     RTNETLINK answers: Invalid argument

     ip xfrm state add src 192.0.2.1 dst 192.0.2.2 proto esp spi 1 \
                reqid 1 mode tunnel aead 'rfc4106(gcm(aes))' \
                0x1111111111111111111111111111111111111111 96 if_id 0
     RTNETLINK answers: Invalid argument

    v1->v2 change:
     - add Fixes: tag

    Fixes: 9f8550e4bd9d ("xfrm: fix disable_xfrm sysctl when used on xfrm interfaces")
    Signed-off-by: Antony Antony <antony.antony@secunet.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 321fc37d299cb9161f145842cc3426b501af8cfd
Author: Antony Antony <antony.antony@secunet.com>
Date:   Sun Dec 12 11:34:30 2021 +0100

    xfrm: interface with if_id 0 should return error

    [ Upstream commit 8dce43919566f06e865f7e8949f5c10d8c2493f5 ]

    xfrm interface if_id = 0 would cause xfrm policy lookup errors since
    Commit 9f8550e4bd9d.

    Now explicitly fail to create an xfrm interface when if_id = 0

    With this commit:
     ip link add ipsec0  type xfrm dev lo  if_id 0
     Error: if_id must be non zero.

    v1->v2 change:
     - add Fixes: tag

    Fixes: 9f8550e4bd9d ("xfrm: fix disable_xfrm sysctl when used on xfrm interfaces")
    Signed-off-by: Antony Antony <antony.antony@secunet.com>
    Reviewed-by: Eyal Birger <eyal.birger@gmail.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 02019dbc7923db007f228886c198664e9107fec6
Author: Jernej Skrabec <jernej.skrabec@gmail.com>
Date:   Mon Nov 29 19:26:25 2021 +0100

    media: hantro: Fix probe func error path

    [ Upstream commit 37af43b250fda6162005d47bf7c959c70d52b107 ]

    If clocks for some reason couldn't be enabled, probe function returns
    immediately, without disabling PM. This obviously leaves PM ref counters
    unbalanced.

    Fix that by jumping to appropriate error path, so effects of PM functions
    are reversed.

    Fixes: 775fec69008d ("media: add Rockchip VPU JPEG encoder driver")
    Signed-off-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Acked-by: Andrzej Pietrasiewicz <andrzej.p@collabora.com>
```

```
    Reviewed-by: Ezequiel Garcia <ezequiel@vanguardiasur.com.ar>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa7e259fe6efe974b29096609d9a378d4c2721a1
Author: Robin Murphy <robin.murphy@arm.com>
Date:   Fri Dec 10 17:54:44 2021 +0000

    drm/tegra: vic: Fix DMA API misuse

    [ Upstream commit 5566174cb10a5167d59b0793871cab7990b149b8 ]

    Upon failure, dma_alloc_coherent() returns NULL. If that does happen,
    passing some uninitialised stack contents to dma_mapping_error() - which
    belongs to a different API in the first place - has precious little
    chance of detecting it.

    Also include the correct header, because the fragile transitive
    inclusion currently providing it is going to break soon.

    Fixes: 20e7dce255e9 ("drm/tegra: Remove memory allocation from Falcon library")
    CC: Thierry Reding <thierry.reding@gmail.com>
    CC: Mikko Perttunen <mperttunen@nvidia.com>
    CC: dri-devel@lists.freedesktop.org
    Signed-off-by: Robin Murphy <robin.murphy@arm.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ebfe49fbd4377783323a6fa196f72f0d7a3c9360
Author: Thierry Reding <treding@nvidia.com>
Date:   Thu Jul 8 16:37:36 2021 +0200

    drm/tegra: gr2d: Explicitly control module reset

    [ Upstream commit 271fca025a6d43f1c18a48543c5aaf31a31e4694 ]

    As of commit 4782c0a5dd88 ("clk: tegra: Don't deassert reset on enabling
    clocks"), module resets are no longer automatically deasserted when the
    module clock is enabled. To make sure that the gr2d module continues to
    work, we need to explicitly control the module reset.

    Fixes: 4782c0a5dd88 ("clk: tegra: Don't deassert reset on enabling clocks")
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db1ae6daa750a76a71fea93f71f206b788855bd0
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Sep 27 11:36:59 2021 +0200

    gpu: host1x: select CONFIG_DMA_SHARED_BUFFER

    [ Upstream commit 6c7a388b62366f0de9936db3c1921d7f4e0011bc ]

    Linking fails when dma-buf is disabled:

    ld.lld: error: undefined symbol: dma_fence_release
    >>> referenced by fence.c
    >>>               gpu/host1x/fence.o:(host1x_syncpt_fence_enable_signaling) in archive
drivers/built-in.a
    >>> referenced by fence.c
    >>>               gpu/host1x/fence.o:(host1x_fence_signal) in archive drivers/built-in.a
    >>> referenced by fence.c
    >>>               gpu/host1x/fence.o:(do_fence_timeout) in archive drivers/built-in.a
```

```
    Fixes: 687db2207b1b ("gpu: host1x: Add DMA fence implementation")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Mikko Perttunen <mperttunen@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df259aaed0b0d2abb099c149f11ac4f08e8a6fa4
Author: Stephen Boyd <swboyd@chromium.org>
Date:   Tue Dec 14 16:25:29 2021 -0800

    drm/bridge: ti-sn65dsi86: Set max register for regmap

    [ Upstream commit 0b665d4af35837f0a0ae63135b84a3c187c1db3b ]

    Set the maximum register to 0xff so we can dump the registers for this
    device in debugfs.

    Fixes: a095f15c00e2 ("drm/bridge: add support for sn65dsi86 bridge driver")
    Cc: Rob Clark <robdclark@chromium.org>
    Cc: Douglas Anderson <dianders@chromium.org>
    Cc: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Reviewed-by: Robert Foss <robert.foss@linaro.org>
    Signed-off-by: Robert Foss <robert.foss@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211215002529.382383-1-
swboyd@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ebed90bd9afbe856bdf9a0f2f986096cb321a76c
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Thu Dec 2 01:26:27 2021 +0300

    drm/msm/dpu: fix safe status debugfs file

    [ Upstream commit f31b0e24d31e18b4503eeaf0032baeacc0beaff6 ]

    Make safe_status debugfs fs file actually return safe status rather than
    danger status data.

    Fixes: 25fdd5933e4c ("drm/msm: Add SDM845 DPU support")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Link: https://lore.kernel.org/r/20211201222633.2476780-3-dmitry.baryshkov@linaro.org
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 260bbe9dee307283a6704b21b1a3ee4b8a35b000
Author: Baruch Siach <baruch@tkos.co.il>
Date:   Tue Dec 7 09:27:10 2021 +0200

    arm64: dts: qcom: ipq6018: Fix gpio-ranges property

    [ Upstream commit 72cb4c48a46a7cfa58eb5842c0d3672ddd5bd9ad ]

    There must be three parameters in gpio-ranges property. Fixes this not
    very helpful error message:

      OF: /soc/pinctrl@1000000: (null) = 3 found 3

    Fixes: 1e8277854b49 ("arm64: dts: Add ipq6018 SoC and CP01 board support")
    Cc: Sricharan R <sricharan@codeaurora.org>
    Signed-off-by: Baruch Siach <baruch@tkos.co.il>
    Tested-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
```

```
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link:
https://lore.kernel.org/r/8a744cfd96aff5754bfdcf7298d208ddca5b319a.1638862030.git.baruch@tko
s.co.il
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f659dbd1cfc27aee71293fcdd4fce739bdf7233
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Thu Dec 9 17:53:41 2021 +0000

    arm64: dts: qcom: c630: Fix soundcard setup

    [ Upstream commit c02b360ca67ebeb9de07b47b2fe53f964c2561d1 ]

    Currently Soundcard has 1 rx device for headset and SoundWire Speaker Playback.

    This setup has issues, ex if we try to play on headset the audio stream is
    also sent to SoundWire Speakers and we will hear sound in both headsets and speakers.

    Make a separate device for Speakers and Headset so that the streams are
    different and handled properly.

    Fixes: 45021d35fcb2 ("arm64: dts: qcom: c630: Enable audio support")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Tested-by: Steev Klimaszewski <steev@kali.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211209175342.20386-2-srinivas.kandagatla@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a8e65665491d688d4d34dd2601c69d54698a109
Author: Kurt Kanzenbach <kurt@linutronix.de>
Date:   Tue Dec 14 14:45:08 2021 +0100

    net: dsa: hellcreek: Add missing PTP via UDP rules

    [ Upstream commit 6cf01e451599da630ff1af529d61c5e4db4550ab ]

    The switch supports PTP for UDP transport too. Therefore, add the missing static
    FDB entries to ensure correct forwarding of these packets.

    Fixes: ddd56dfe52c9 ("net: dsa: hellcreek: Add PTP clock support")
    Signed-off-by: Kurt Kanzenbach <kurt@linutronix.de>
    Acked-by: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7f82c5da785a970bd35d83fa483fc6909d96c137
Author: Kurt Kanzenbach <kurt@linutronix.de>
Date:   Tue Dec 14 14:45:07 2021 +0100

    net: dsa: hellcreek: Allow PTP P2P measurements on blocked ports

    [ Upstream commit cad1798d2d0811ded37d1e946c6796102e58013b ]

    Allow PTP peer delay measurements on blocked ports by STP. In case of topology
    changes the PTP stack can directly start with the correct delays.

    Fixes: ddd56dfe52c9 ("net: dsa: hellcreek: Add PTP clock support")
    Signed-off-by: Kurt Kanzenbach <kurt@linutronix.de>
    Reviewed-by: Vladimir Oltean <olteanv@gmail.com>
    Acked-by: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9f9007ea07fcece097210684046962758 8a6051
```

```
Author: Kurt Kanzenbach <kurt@linutronix.de>
Date:   Tue Dec 14 14:45:06 2021 +0100

    net: dsa: hellcreek: Add STP forwarding rule

    [ Upstream commit b7ade35eb53a2455f737a623c24e4b24455b2271 ]

    Treat STP as management traffic. STP traffic is designated for the CPU port
    only. In addition, STP traffic has to pass blocked ports.

    Fixes: e4b27ebc780f ("net: dsa: Add DSA driver for Hirschmann Hellcreek switches")
    Signed-off-by: Kurt Kanzenbach <kurt@linutronix.de>
    Reviewed-by: Vladimir Oltean <olteanv@gmail.com>
    Acked-by: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit af1629fe51121e81ff3179a5c12ce9167a9fe912
Author: Kurt Kanzenbach <kurt@linutronix.de>
Date:   Tue Dec 14 14:45:05 2021 +0100

    net: dsa: hellcreek: Fix insertion of static FDB entries

    [ Upstream commit 4db4c3ea56978086ca367a355e440de17d534827 ]

    The insertion of static FDB entries ignores the pass_blocked bit. That bit is
    evaluated with regards to STP. Add the missing functionality.

    Fixes: e4b27ebc780f ("net: dsa: Add DSA driver for Hirschmann Hellcreek switches")
    Signed-off-by: Kurt Kanzenbach <kurt@linutronix.de>
    Reviewed-by: Vladimir Oltean <olteanv@gmail.com>
    Acked-by: Richard Cochran <richardcochran@gmail.com>
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94412bcafa3f7075c61df0c51657c4c4f96e17f9
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Mon Dec 13 11:53:07 2021 +0200

    ath11k: Fix a NULL pointer dereference in ath11k_mac_op_hw_scan()

    [ Upstream commit eccd25136386a04ebf46a64f3a34e8e0fab6d9e1 ]

    In ath11k_mac_op_hw_scan(), the return value of kzalloc() is directly
    used in memcpy(), which may lead to a NULL pointer dereference on
    failure of kzalloc().

    Fix this bug by adding a check of arg.extraie.ptr.

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_ATH11K=m show no new warnings, and our static
    analyzer no longer warns about this code.

    Fixes: d5c65159f289 ("ath11k: driver for Qualcomm IEEE 802.11ax devices")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
```

```
    Link: https://lore.kernel.org/r/20211202155348.71315-1-zhou1615@umn.edu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d9c6e0e506a702ed9db7aa7441bc54dd6cb20e9
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Mon Dec 6 03:22:01 2021 +0100

    media: coda/imx-vdoa: Handle dma_set_coherent_mask error codes

    [ Upstream commit 43f0633f89947df57fe0b5025bdd741768007708 ]

    The return value of dma_set_coherent_mask() is not always 0.
    To catch the exception in case that dma is not support the mask.

    Link: https://lore.kernel.org/linux-media/20211206022201.1639460-1-jiasheng@iscas.ac.cn
    Fixes: b0444f18e0b1 ("[media] coda: add i.MX6 VDOA driver")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f55d24ba87a8cc64ca63fab59c695b6daa1c328
Author: Wang Hai <wanghai38@huawei.com>
Date:   Tue Oct 26 13:23:48 2021 +0200

    media: msi001: fix possible null-ptr-deref in msi001_probe()

    [ Upstream commit 3d5831a40d3464eea158180eb12cbd81c5edfb6a ]

    I got a null-ptr-deref report:

    BUG: kernel NULL pointer dereference, address: 0000000000000060
    ...
    RIP: 0010:v4l2_ctrl_auto_cluster+0x57/0x270
    ...
    Call Trace:
     msi001_probe+0x13b/0x24b [msi001]
     spi_probe+0xeb/0x130
    ...
     do_syscall_64+0x35/0xb0

    In msi001_probe(), if the creation of control for bandwidth_auto
    fails, there will be a null-ptr-deref issue when it is used in
    v4l2_ctrl_auto_cluster().

    Check dev->hdl.error before v4l2_ctrl_auto_cluster() to fix this bug.

    Link: https://lore.kernel.org/linux-media/20211026112348.2878040-1-wanghai38@huawei.com
    Fixes: 93203dd6c7c4 ("[media] msi001: Mirics MSi001 silicon tuner driver")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Hai <wanghai38@huawei.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 416c6fa2ec32a8aa887ace2e1a4ef8cbf15156c3
Author: Anton Vasilyev <vasilyev@ispras.ru>
Date:   Thu Aug 22 12:41:47 2019 +0200

    media: dw2102: Fix use after free

    [ Upstream commit 589a9f0eb799f77de2c09583bf5bad221fa5d685 ]

    dvb_usb_device_init stores parts of properties at d->props
    and d->desc and uses it on dvb_usb_device_exit.
    Free of properties on module probe leads to use after free.
    Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=204597
```

```
    The patch makes properties static instead of allocated on heap to prevent
    memleak and use after free.
    Also fixes s421_properties.devices initialization to have 2 element
    instead of 6 copied from p7500_properties.

    [mchehab: fix function call alignments]
    Link: https://lore.kernel.org/linux-media/20190822104147.4420-1-vasilyev@ispras.ru
    Signed-off-by: Anton Vasilyev <vasilyev@ispras.ru>
    Fixes: 299c7007e936 ("media: dw2102: Fix memleak on sequence of probes")
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3dce8ccca7ac468e4cb6d8ccb24014d1b89f74b
Author: Robin Murphy <robin.murphy@arm.com>
Date:   Fri Dec 3 11:44:50 2021 +0000

    perf/arm-cmn: Fix CPU hotplug unregistration

    [ Upstream commit 56c7c6eaf3eb8ac1ec40d56096c0f2b27250da5f ]

    Attempting to migrate the PMU context after we've unregistered the PMU
    device, or especially if we never successfully registered it in the
    first place, is a woefully bad idea. It's also fundamentally pointless
    anyway. Make sure to unregister an instance from the hotplug handler
    *without* invoking the teardown callback.

    Fixes: 0ba64770a2f2 ("perf: Add Arm CMN-600 PMU driver")
    Signed-off-by: Robin Murphy <robin.murphy@arm.com>
    Link:
https://lore.kernel.org/r/2c221d745544774e4b07583b65b5d4d94f7e0fe4.1638530442.git.robin.murp
hy@arm.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc8bd9347528abb13b2a99b317b519f7d2389831
Author: Christian Lamparter <chunkeey@gmail.com>
Date:   Mon Dec 6 01:43:34 2021 +0100

    ARM: dts: gemini: NAS4220-B: fis-index-block with 128 KiB sectors

    [ Upstream commit 4754eab7e5a78bdefe7a960c5c260c95ebbb5fa6 ]

    Steven Maddox reported in the OpenWrt bugzilla, that his
    RaidSonic IB-NAS4220-B was no longer booting with the new
    OpenWrt 21.02 (uses linux 5.10's device-tree). However, it was
    working with the previous OpenWrt 19.07 series (uses 4.14).

    |[    5.548038] No RedBoot partition table detected in 30000000.flash
    |[    5.618553] Searching for RedBoot partition table in 30000000.flash at offset 0x0
    |[    5.739093] No RedBoot partition table detected in 30000000.flash
    |...
    |[    7.039504] Waiting for root device /dev/mtdblock3...

    The provided bootlog shows that the RedBoot partition parser was
    looking for the partition table "at offset 0x0". Which is strange
    since the comment in the device-tree says it should be at 0xfe0000.

    Further digging on the internet led to a review site that took
    some useful PCB pictures of their review unit back in February 2009.
    Their picture shows a Spansion S29GL128N11TFI01 flash chip.

    >From Spansion's Datasheet:
    "S29GL128N: One hundred twenty-eight 64 Kword (128 Kbyte) sectors"
    Steven also provided a "cat /sys/class/mtd/mtd0/erasesize" from his
```

```
      unit: "131072".

      With the 128 KiB Sector/Erasesize in mind. This patch changes the
      fis-index-block property to (0xfe0000 / 0x20000) = 0x7f.

      Fixes: b5a923f8c739 ("ARM: dts: gemini: Switch to redboot partition parsing")
      Reported-by: Steven Maddox <s.maddox@lantizia.me.uk>
      Signed-off-by: Christian Lamparter <chunkeey@gmail.com>
      Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
      Tested-by: Steven Maddox <s.maddox@lantizia.me.uk>
      Link: https://lore.kernel.org/r/20211206004334.4169408-1-linus.walleij@linaro.org'
      Bugzilla: https://bugs.openwrt.org/index.php?do=details&task_id=4137
      Signed-off-by: Arnd Bergmann <arnd@arndb.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit d7e9cd51d60df1c1999c1c05a253353e2053c4cd
  Author: Daniel Latypov <dlatypov@google.com>
  Date:   Fri Oct 8 16:24:21 2021 -0700

      kunit: tool: fix --json output for skipped tests

      [ Upstream commit 9a6bb30a8830bb868b09629f0b9ad5d2b5fbb2f9 ]

      Currently, KUnit will report SKIPPED tests as having failed if one uses
      --json.

      Add the missing if statement to set the appropriate status ("SKIP").
      See https://api.kernelci.org/schema-test-case.html:
        "status": {
            "type": "string",
            "description": "The status of the execution of this test case",
            "enum": ["PASS", "FAIL", "SKIP", "ERROR"],
            "default": "PASS"
        },
      with this, we now can properly produce all four of the statuses.

      Fixes: 5acaf6031f53 ("kunit: tool: Support skipped tests in kunit_tool")
      Signed-off-by: Daniel Latypov <dlatypov@google.com>
      Reviewed-by: David Gow <davidgow@google.com>
      Reviewed-by: Brendan Higgins <brendanhiggins@google.com>
      Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 959bf0b595fe974f8a3c8bd9401a49d2e96e5283
  Author: Hector Martin <marcan@marcan.st>
  Date:   Sat Dec 11 02:05:34 2021 +0900

      spi: Fix incorrect cs_setup delay handling

      [ Upstream commit 95c07247399536f83b89dc60cfe7b279d17e69f6 ]

      Move the cs_setup delay to the end of spi_set_cs.

      From include/linux/spi/spi.h:

       * @cs_setup: delay to be introduced by the controller after CS is
         asserted

      The cs_setup delay needs to happen *after* CS is asserted, that is, at
      the end of spi_set_cs, not at the beginning. Otherwise we're just
      delaying before the SPI transaction starts at all, which isn't very
      useful.

      No drivers use this right now, but that is likely to change soon with an
      upcoming Apple SPI HID transport driver.
```

```
    Fixes: 25093bdeb6bc ("spi: implement SW control for CS times")
    Signed-off-by: Hector Martin <marcan@marcan.st>
    Link: https://lore.kernel.org/r/20211210170534.177139-1-marcan@marcan.st
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfb81bda4f76129fada681659303c2e5d41d174e
Author: Rameshkumar Sundaram <quic_ramess@quicinc.com>
Date:   Thu Dec 9 23:07:01 2021 +0530

    ath11k: Fix deleting uninitialized kernel timer during fragment cache flush

    [ Upstream commit ba53ee7f7f38cf0592b8be1dcdabaf8f7535f8c1 ]

    frag_timer will be created & initialized for stations when
    they associate and will be deleted during every key installation
    while flushing old fragments.

    For AP interface self peer will be created and Group keys
    will be installed for this peer, but there will be no real
    Station entry & hence frag_timer won't be created and
    initialized, deleting such uninitialized kernel timers causes below
    warnings and backtraces printed with CONFIG_DEBUG_OBJECTS_TIMERS
    enabled.

    [ 177.828008] ODEBUG: assert_init not available (active state 0) object type: timer_list
hint: 0x0
    [ 177.836833] WARNING: CPU: 3 PID: 188 at lib/debugobjects.c:508
debug_print_object+0xb0/0xf0
    [ 177.845185] Modules linked in: ath11k_pci ath11k qmi_helpers qrtr_mhi qrtr ns mhi
    [ 177.852679] CPU: 3 PID: 188 Comm: hostapd Not tainted 5.14.0-rc3-32919-g4034139e1838-
dirty #14
    [ 177.865805] pstate: 60000005 (nZCv daif -PAN -UAO -TCO BTYPE=--)
    [ 177.871804] pc : debug_print_object+0xb0/0xf0
    [ 177.876155] lr : debug_print_object+0xb0/0xf0
    [ 177.880505] sp : ffffffc01169b5a0
    [ 177.883810] x29: ffffffc01169b5a0 x28: ffffff80081c2320 x27: ffffff80081c4078
    [ 177.890942] x26: ffffff8003fe8f28 x25: ffffff8003de9890 x24: ffffffc01134d738
    [ 177.898075] x23: ffffffc010948f20 x22: ffffffc010b2d2e0 x21: ffffffc01169b628
    [ 177.905206] x20: ffffffc01134d700 x19: ffffffc010c80d98 x18: 00000000000003f6
    [ 177.912339] x17: 203a657079742074 x16: 63656a626f202930 x15: 0000000000000152
    [ 177.919471] x14: 0000000000000152 x13: 00000000ffffffea x12: ffffffc010d732e0
    [ 177.926603] x11: 0000000000000003 x10: ffffffc010d432a0 x9 : ffffffc010d432f8
    [ 177.933735] x8 : 000000000002ffe8 x7 : c0000000ffffdfff x6 : 0000000000000001
    [ 177.940866] x5 : 0000000000000000 x4 : 0000000000000000 x3 : 00000000ffffffff
    [ 177.947997] x2 : ffffffc010c93240 x1 : ffffff80023624c0 x0 : 0000000000000054
    [ 177.955130] Call trace:
    [ 177.957567] debug_print_object+0xb0/0xf0
    [ 177.961570] debug_object_assert_init+0x124/0x178
    [ 177.966269] try_to_del_timer_sync+0x1c/0x70
    [ 177.970536] del_timer_sync+0x30/0x50
    [ 177.974192] ath11k_peer_frags_flush+0x34/0x68 [ath11k]
    [ 177.979439] ath11k_mac_op_set_key+0x1e4/0x338 [ath11k]
    [ 177.984673] ieee80211_key_enable_hw_accel+0xc8/0x3d0
    [ 177.989722] ieee80211_key_replace+0x360/0x740
    [ 177.994160] ieee80211_key_link+0x16c/0x210
    [ 177.998337] ieee80211_add_key+0x138/0x338
    [ 178.002426] nl80211_new_key+0xfc/0x258
    [ 178.006257] genl_family_rcv_msg_doit.isra.17+0xd8/0x120
    [ 178.011565] genl_rcv_msg+0xd8/0x1c8
    [ 178.015134] netlink_rcv_skb+0x38/0xf8
    [ 178.018877] genl_rcv+0x34/0x48
    [ 178.022012] netlink_unicast+0x174/0x230
    [ 178.025928] netlink_sendmsg+0x188/0x388
```

```
     [ 178.029845] ____sys_sendmsg+0x218/0x250
     [ 178.033763] ___sys_sendmsg+0x68/0x90
     [ 178.037418] __sys_sendmsg+0x44/0x88
     [ 178.040988] __arm64_sys_sendmsg+0x20/0x28
     [ 178.045077] invoke_syscall.constprop.5+0x54/0xe0
     [ 178.049776] do_el0_svc+0x74/0xc0
     [ 178.053084] el0_svc+0x10/0x18
     [ 178.056133] el0t_64_sync_handler+0x88/0xb0
     [ 178.060310] el0t_64_sync+0x148/0x14c
     [ 178.063966] ---[ end trace 8a5cf0bf9d34a058 ]---

     Add changes to not to delete frag timer for peers during
     group key installation.

     Tested on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-01092-QCAHKSWPL_SILICONZ-1

     Fixes: c3944a562102 ("ath11k: Clear the fragment cache during key install")
     Signed-off-by: Rameshkumar Sundaram <quic_ramess@quicinc.com>
     Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
     Link: https://lore.kernel.org/r/1639071421-25078-1-git-send-email-
  quic_ramess@quicinc.com
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 8f485f3fd73ecd3c25c519a6a38b0ea438bbe822
  Author: Alexei Starovoitov <ast@kernel.org>
  Date:   Sat Dec 11 17:16:19 2021 -0800

     libbpf: Fix gen_loader assumption on number of programs.

     [ Upstream commit 259172bb6514758ce3be1610c500b51a9f44212a ]

     libbpf's obj->nr_programs includes static and global functions. That number
     could be higher than the actual number of bpf programs going be loaded by
     gen_loader. Passing larger nr_programs to bpf_gen__init() doesn't hurt. Those
     exra stack slots will stay as zero. bpf_gen__finish() needs to check that
     actual number of progs that gen_loader saw is less than or equal to
     obj->nr_programs.

     Fixes: ba05fd36b851 ("libbpf: Perform map fd cleanup for gen_loader in case of error")
     Signed-off-by: Alexei Starovoitov <ast@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 3476961a9fad8d0882d7d523992fa3cefd8e265c
  Author: Weili Qian <qianweili@huawei.com>
  Date:   Sat Dec 4 18:43:01 2021 +0800

     crypto: hisilicon/qm - fix incorrect return value of hisi_qm_resume()

     [ Upstream commit 3f9dd4c802b96626e869b2d29c8e401dabadd23e ]

     When hisi_qm_resume() returns 0, it indicates that the device has started
     successfully.  If the device fails to start, hisi_qm_resume() needs to
     return the actual error code to the caller instead of 0.

     Fixes: d7ea53395b72 ("crypto: hisilicon - add runtime PM ops")
     Signed-off-by: Weili Qian <qianweili@huawei.com>
     Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 80d2ddd428ffac6dd0cc006832359019ea6faf54
  Author: Herbert Xu <herbert@gondor.apana.org.au>
  Date:   Wed Dec 1 17:30:41 2021 +1100

     crypto: stm32 - Revert broken pm_runtime_resume_and_get changes
```

```
    [ Upstream commit 3d6b661330a7954d8136df98160d525eb04dcd6a ]

    We should not call pm_runtime_resume_and_get where the reference
    count is expected to be incremented unconditionally.  This patch
    reverts these calls to the original unconditional get_sync call.

    Reported-by: Heiner Kallweit <hkallweit1@gmail.com>
    Fixes: 747bf30fd944 ("crypto: stm32/cryp - Fix PM reference leak...")
    Fixes: 1cb3ad701970 ("crypto: stm32/hash - Fix PM reference leak...")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Acked-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97f100d9435f5ff9e18c07a2469cc1f2d3aac2d3
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:55:00 2021 +0100

    crypto: stm32/cryp - fix bugs and crash in tests

    [ Upstream commit 4b898d5cfa4d9a0ad5bc82cb5eafdc092394c6a9 ]

    Extra crypto manager auto test were crashing or failling due
    to 2 reasons:
    - block in a dead loop (dues to issues in cipher end process management)
    - crash due to read/write unmapped memory (this crash was also reported
    when using openssl afalg engine)

    Rework interrupt management, interrupts are masked as soon as they are
    no more used: if input buffer is fully consumed, "Input FIFO not full"
    interrupt is masked and if output buffer is full, "Output FIFO not
    empty" interrupt is masked.
    And crypto request finish when input *and* outpout buffer are fully
    read/write.

    About the crash due to unmapped memory, using scatterwalk_copychunks()
    that will map and copy each block fix the issue.
    Using this api and copying full block will also fix unaligned data
    access, avoid early copy of in/out buffer, and make useless the extra
    alignment constraint.

    Fixes: 9e054ec21ef8 ("crypto: stm32 - Support for STM32 CRYP crypto module")

    Reported-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c88e4241bb1dda01dbeaaf57b9435dd0b7814b33
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:54:59 2021 +0100

    crypto: stm32/cryp - fix lrw chaining mode

    [ Upstream commit fa97dc2d48b476ea98199d808d3248d285987e99 ]

    This fixes the lrw autotest if lrw uses the CRYP as the AES block cipher
    provider (as ecb(aes)). At end of request, CRYP should not update the IV
    in case of ECB chaining mode. Indeed the ECB chaining mode never uses
    the IV, but the software LRW chaining mode uses the IV field as
    a counter and due to the (unexpected) update done by CRYP while the AES
    block process, the counter get a wrong value when the IV overflow.

    Fixes: 5f49f18d27cd ("crypto: stm32/cryp - update to return iv_out")
```

```
    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9b97b1ab0d4588f25d82b494ed6f28b3bec240b
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:54:58 2021 +0100

    crypto: stm32/cryp - fix double pm exit

    [ Upstream commit 6c12e742785bf9333faf60bfb96575bdd763448e ]

    Delete extraneous lines in probe error handling code: pm was
    disabled twice.

    Fixes: 65f9aa36ee47 ("crypto: stm32/cryp - Add power management support")

    Reported-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48aa7e96d974c418a3d030facec0db8b337f3935
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:54:57 2021 +0100

    crypto: stm32/cryp - check early input data

    [ Upstream commit 39e6e699c7fb92bdb2617b596ca4a4ea35c5d2a7 ]

    Some auto tests failed because driver wasn't returning the expected
    error with some input size/iv value/tag size.
    Now:
     Return 0 early for empty buffer. (We don't need to start the engine for
     an empty input buffer).
     Accept any valid authsize for gcm(aes).
     Return -EINVAL if iv for ccm(aes) is invalid.
     Return -EINVAL if buffer size is a not a multiple of algorithm block size.

    Fixes: 9e054ec21ef8 ("crypto: stm32 - Support for STM32 CRYP crypto module")

    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89904a3ee053a8c68b0d87ab9768a5cbae3f1127
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:54:56 2021 +0100

    crypto: stm32/cryp - fix xts and race condition in crypto_engine requests

    [ Upstream commit d703c7a994ee34b7fa89baf21631fca0aa9f17fc ]

    Don't erase key:
    If key is erased before the crypto_finalize_.*_request() call, some
    pending process will run with a key={ 0 }.
    Moreover if the key is reset at end of request, it breaks xts chaining
    mode, as for last xts block (in case input len is not a multiple of
    block) a new AES request is started without calling again set_key().

    Fixes: 9e054ec21ef8 ("crypto: stm32 - Support for STM32 CRYP crypto module")

    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 28cce5a3d884bf91af7884468425b3943d8e3790
Author: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
Date:   Tue Nov 30 08:54:55 2021 +0100

    crypto: stm32/cryp - fix CTR counter carry

    [ Upstream commit 41c76690b0990efacd15d35cfb4e77318cd80ebb ]

    STM32 CRYP hardware doesn't manage CTR counter bigger than max U32, as
    a workaround, at each block the current IV is saved, if the saved IV
    lower u32 is 0xFFFFFFFF, the full IV is manually incremented, and set
    in hardware.
    Fixes: bbb2832620ac ("crypto: stm32 - Fix sparse warnings")

    Signed-off-by: Nicolas Toromanoff <nicolas.toromanoff@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f16ec52953f931d6c40761358d6c1bb3f5dffa4e
Author: Jakub Kicinski <kuba@kernel.org>
Date:   Wed Nov 24 14:39:16 2021 -0800

    selftests: harness: avoid false negatives if test has no ASSERTs

    [ Upstream commit 3abedf4646fdc0036fcb8ebbc3b600667167fafe ]

    Test can fail either immediately when ASSERT() failed or at the
    end if one or more EXPECT() was not met. The exact return code
    is decided based on the number of successful ASSERT()s.

    If test has no ASSERT()s, however, the return code will be 0,
    as if the test did not fail. Start counting ASSERT()s from 1.

    Fixes: 369130b63178 ("selftests: Enhance kselftest_harness.h to print which assert
failed")
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66a82e499e708c087e185b834a1ba17105f26e67
Author: Anders Roxell <anders.roxell@linaro.org>
Date:   Wed Nov 3 21:13:50 2021 +0100

    selftests: clone3: clone3: add case CLONE3_ARGS_NO_TEST

    [ Upstream commit a531b0c23c0fc68ad758cc31a74cf612a4dafeb0 ]

    Building selftests/clone3 with clang warns about enumeration not handled
    in switch case:

    clone3.c:54:10: warning: enumeration value 'CLONE3_ARGS_NO_TEST' not handled in switch
[-Wswitch]
            switch (test_mode) {
                    ^

    Add the missing switch case with a comment.

    Fixes: 17a810699c18 ("selftests: add tests for clone3()")
    Signed-off-by: Anders Roxell <anders.roxell@linaro.org>
    Acked-by: Christian Brauner <christian.brauner@ubuntu.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e81e03ee75bb0647d108e7d07c5852e697ae4cc7
```

```
Author: Shuyi Cheng <chengshuyi@linux.alibaba.com>
Date:   Fri Dec 10 17:39:57 2021 +0800

    libbpf: Add "bool skipped" to struct bpf_map

    [ Upstream commit 229fae38d0fc0d6ff58d57cbeb1432da55e58d4f ]

    Fix error: "failed to pin map: Bad file descriptor, path:
    /sys/fs/bpf/_rodata_str1_1."

    In the old kernel, the global data map will not be created, see [0]. So
    we should skip the pinning of the global data map to avoid
    bpf_object__pin_maps returning error. Therefore, when the map is not
    created, we mark “map->skipped" as true and then check during relocation
    and during pinning.

    Fixes: 16e0c35c6f7a ("libbpf: Load global data maps lazily on legacy kernels")
    Signed-off-by: Shuyi Cheng <chengshuyi@linux.alibaba.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 17bf2fbf0afbbf09cd6ddc52c0260f0f964fb672
Author: Kees Cook <keescook@chromium.org>
Date:   Wed Dec 8 20:34:56 2021 -0800

    x86/uaccess: Move variable into switch case statement

    [ Upstream commit 61646ca83d3889696f2772edaff122dd96a2935e ]

    When building with automatic stack variable initialization, GCC 12
    complains about variables defined outside of switch case statements.
    Move the variable into the case that uses it, which silences the warning:

    ./arch/x86/include/asm/uaccess.h:317:23: warning: statement will never be executed [-
Wswitch-unreachable]
      317 |         unsigned char x_u8__; \
          |                       ^~~~~~

    Fixes: 865c50e1d279 ("x86/uaccess: utilize CONFIG_CC_HAS_ASM_GOTO_OUTPUT")
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Dave Hansen <dave.hansen@linux.intel.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lkml.kernel.org/r/20211209043456.1377875-1-keescook@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2825f30c449a2476460c4df7195449cd19f4f504
Author: Eric Dumazet <edumazet@google.com>
Date:   Wed Dec 8 12:20:19 2021 -0800

    xfrm: fix a small bug in xfrm_sa_len()

    [ Upstream commit 7770a39d7c63faec6c4f33666d49a8cb664d0482 ]

    copy_user_offload() will actually push a struct struct xfrm_user_offload,
    which is different than (struct xfrm_state *)->xso
    (struct xfrm_state_offload)

    Fixes: d77e38e612a01 ("xfrm: Add an IPsec hardware offloading API")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Steffen Klassert <steffen.klassert@secunet.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b30e4c2788ac718bdbb5f729aea9d0581958b71c
Author: Brian Norris <briannorris@chromium.org>
```

```
Date:   Mon Nov 29 16:47:34 2021 -0800

    mwifiex: Fix possible ABBA deadlock

    [ Upstream commit 1b8bb8919ef81bfc8873d223b9361f1685f2106d ]

    Quoting Jia-Ju Bai <baijiaju1990@gmail.com>:

      mwifiex_dequeue_tx_packet()
         spin_lock_bh(&priv->wmm.ra_list_spinlock); --> Line 1432 (Lock A)
         mwifiex_send_addba()
           spin_lock_bh(&priv->sta_list_spinlock); --> Line 608 (Lock B)

      mwifiex_process_sta_tx_pause()
         spin_lock_bh(&priv->sta_list_spinlock); --> Line 398 (Lock B)
         mwifiex_update_ralist_tx_pause()
           spin_lock_bh(&priv->wmm.ra_list_spinlock); --> Line 941 (Lock A)

    Similar report for mwifiex_process_uap_tx_pause().

    While the locking expectations in this driver are a bit unclear, the
    Fixed commit only intended to protect the sta_ptr, so we can drop the
    lock as soon as we're done with it.

    IIUC, this deadlock cannot actually happen, because command event
    processing (which calls mwifiex_process_sta_tx_pause()) is
    sequentialized with TX packet processing (e.g.,
    mwifiex_dequeue_tx_packet()) via the main loop (mwifiex_main_process()).
    But it's good not to leave this potential issue lurking.

    Fixes: f0f7c2275fb9 ("mwifiex: minor cleanups w/ sta_list_spinlock in cfg80211.c")
    Cc: Douglas Anderson <dianders@chromium.org>
    Reported-by: TOTE Robot <oslab@tsinghua.edu.cn>
    Link: https://lore.kernel.org/linux-wireless/0e495b14-efbb-e0da-37bd-
af6bd677ee2c@gmail.com/
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/YaV0pllJ5p/EuUat@google.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6295d9568992bf5cd31b6b584a733737e54387b4
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Thu Nov 25 21:01:14 2021 +0300

    drm/msm/dsi: fix initialization in the bonded DSI case

    [ Upstream commit 92cb1bedde9dba78d802fe2510949743a2581aed ]

    Commit 739b4e7756d3 ("drm/msm/dsi: Fix an error code in
    msm_dsi_modeset_init()") changed msm_dsi_modeset_init() to return an
    error code in case msm_dsi_manager_validate_current_config() returns
    false. However this is not an error case, but a slave DSI of the bonded
    DSI link. In this case msm_dsi_modeset_init() should return 0, but just
    skip connector and bridge initialization.

    To reduce possible confusion, drop the
    msm_dsi_manager_validate_current_config() function, and specif 'bonded
    && !master' condition directly in the msm_dsi_modeset_init().

    Fixes: 739b4e7756d3 ("drm/msm/dsi: Fix an error code in msm_dsi_modeset_init()")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Link: https://lore.kernel.org/r/20211125180114.561278-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Rob Clark <robdclark@chromium.org>
```

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

commit c828c5f1e535acf797a3b46ff2e69111414fff43
Author: Loic Poulain <loic.poulain@linaro.org>
Date:   Thu Dec 2 10:02:12 2021 +0100

    wcn36xx: Fix max channels retrieval

    [ Upstream commit 09cab4308bf9b8076ee4a3c56015daf9ef9cb23e ]

    Kernel test robot reported:drivers/net/wireless/ath/wcn36xx/smd.c:943:33:
       sparse: sparse: cast truncates bits from constant value (780 becomes 80)

    The 'channels' field is not a simple u8 array but an array of
    channel_params. Using sizeof for retrieving the max number of
    channels is then wrong.

    In practice, it was not an issue, because the sizeof returned
    value is 780, which is truncated in min_t (u8) to 80, which is
    the value we expect...

    Fix that properly using ARRAY_SIZE instead of sizeof.

    Fixes: d707f812bb05 ("wcn36xx: Channel list update before hardware scan")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/1638435732-14657-1-git-send-email-
loic.poulain@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7f44d3c70e94f990defc2b30265e60755b95d64
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Tue Nov 30 17:21:08 2021 +0100

    rcu/exp: Mark current CPU as exp-QS in IPI loop second pass

    [ Upstream commit 81f6d49cce2d2fe507e3fddcc4a6db021d9c2e7b ]

    Expedited RCU grace periods invoke sync_rcu_exp_select_node_cpus(), which
    takes two passes over the leaf rcu_node structure's CPUs.  The first
    pass gathers up the current CPU and CPUs that are in dynticks idle mode.
    The workqueue will report a quiescent state on their behalf later.
    The second pass sends IPIs to the rest of the CPUs, but excludes the
    current CPU, incorrectly assuming it has been included in the first
    pass's list of CPUs.

    Unfortunately the current CPU may have changed between the first and
    second pass, due to the fact that the various rcu_node structures'
    ->lock fields have been dropped, thus momentarily enabling preemption.
    This means that if the second pass's CPU was not on the first pass's
    list, it will be ignored completely.  There will be no IPI sent to
    it, and there will be no reporting of quiescent states on its behalf.
    Unfortunately, the expedited grace period will nevertheless be waiting
    for that CPU to report a quiescent state, but with that CPU having no
    reason to believe that such a report is needed.

    The result will be an expedited grace period stall.

    Fix this by no longer excluding the current CPU from consideration during
    the second pass.

    Fixes: b9ad4d6ed18e ("rcu: Avoid self-IPI in sync_rcu_exp_select_node_cpus()")
    Reviewed-by: Neeraj Upadhyay <quic_neeraju@quicinc.com>
    Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
```

```
      Cc: Uladzislau Rezki <urezki@gmail.com>
      Cc: Neeraj Upadhyay <quic_neeraju@quicinc.com>
      Cc: Boqun Feng <boqun.feng@gmail.com>
      Cc: Josh Triplett <josh@joshtriplett.org>
      Cc: Joel Fernandes <joel@joelfernandes.org>
      Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit a664e575d256d20b67f910a113ccb216f05a87e2
  Author: Jackie Liu <liuyun01@kylinos.cn>
  Date:   Wed Nov 10 15:09:49 2021 +0800

      drm/msm/dp: displayPort driver need algorithm rational

      [ Upstream commit 53d22794711ad630f40d59dd726bd260d77d585f ]

      Let's select RATIONAL with dp driver. avoid like:

      [...]
      x86_64-linux-gnu-ld: drivers/gpu/drm/msm/dp/dp_catalog.o: in function
  `dp_catalog_ctrl_config_msa':
      dp_catalog.c:(.text+0x57e): undefined reference to `rational_best_approximation'

      Fixes: c943b4948b58 ("drm/msm/dp: add displayPort driver support")
      Reported-by: kernelbot <kernel-bot@kylinos.cn>
      Signed-off-by: Jackie Liu <liuyun01@kylinos.cn>
      Link: https://lore.kernel.org/r/20211110070950.3355597-2-liu.yun@linux.dev
      Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
      Signed-off-by: Rob Clark <robdclark@chromium.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 91628c62b4430424d3b7ee6e921ad4edc03cb113
  Author: Rob Clark <robdclark@chromium.org>
  Date:   Mon Nov 29 10:23:44 2021 -0800

      drm/msm/gpu: Don't allow zero fence_id

      [ Upstream commit ca3ffcbeb0c866d9b0cb38eaa2bd4416597b5966 ]

      Elsewhere we treat zero as "no fence" and __msm_gem_submit_destroy()
      skips removal from fence_idr.  We could alternately change this to use
      negative values for "no fence" but I think it is more clear to not allow
      zero as a valid fence_id.

      Signed-off-by: Rob Clark <robdclark@chromium.org>
      Fixes: a61acbbe9cf8 ("drm/msm: Track "seqno" fences by idr")
      Link: https://lore.kernel.org/r/20211129182344.292609-1-robdclark@gmail.com
      Signed-off-by: Rob Clark <robdclark@chromium.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 2491d93334b26043dbbbabb9d51d57aa952c2126
  Author: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
  Date:   Tue Dec 7 09:46:39 2021 -0500

      drm/amd/display: Fix out of bounds access on DNC31 stream encoder regs

      [ Upstream commit d374d3b493215d637b9e7be12a93f22caf4c1f97 ]

      [Why]
      During dcn31_stream_encoder_create, if PHYC/D get remapped to F/G on B0
      then we'll index 5 or 6 into a array of length 5 - leading to an
      access violation on some configs during device creation.

      [How]
      Software won't be touching PHYF/PHYG directly, so just extend the
```

```
    array to cover all possible engine IDs.

    Even if it does by try to access one of these registers by accident
    the offset will be 0 and we'll get a warning during the access.

    Fixes: 2fe9a0e1173f ("drm/amd/display: Fix DCN3 B0 DP Alt Mapping")
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Nicholas Kazlauskas <nicholas.kazlauskas@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit efd6004a0ffb1f54b80e871d6623b8892e0dc650
Author: Wayne Lin <Wayne.Lin@amd.com>
Date:   Fri Nov 26 15:19:57 2021 +0800

    drm/amd/display: Fix bug in debugfs crc_win_update entry

    [ Upstream commit 4bef85d4c9491415b7931407b07f24841c1e0390 ]

    [Why]
    crc_rd_wrk shouldn't be null in crc_win_update_set(). Current programming
    logic is inconsistent in crc_win_update_set().

    [How]
    Initially, return if crc_rd_wrk is NULL. Later on, we can use member of
    crc_rd_wrk safely.

    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Fixes: 9a65df193108 ("drm/amd/display: Use PSP TA to read out crc")

    Reviewed-by: Rodrigo Siqueira <Rodrigo.Siqueira@amd.com>
    Acked-by: Pavle Kotarac <Pavle.Kotarac@amd.com>
    Signed-off-by: Wayne Lin <Wayne.Lin@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e37347305579cba3bc2e8d0443c44620b90655e4
Author: Mark Chen <mark-yw.chen@mediatek.com>
Date:   Tue Dec 7 01:33:42 2021 +0800

    Bluetooth: btusb: Handle download_firmware failure cases

    [ Upstream commit 00c0ee9850b7b0cb7c40b8daba806ae2245e59d4 ]

    For Mediatek chipset, it can not enabled if there are something wrong
    in btmtk_setup_firmware_79xx(). Thus, the process must be terminated
    and returned error code.

    Fixes: fc342c4dc4087 ("Bluetooth: btusb: Add protocol support for MediaTek MT7921U USB
devices")
    Co-developed-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Mark Chen <mark-yw.chen@mediatek.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff030b361629ca953541d074fa6cf9e765c65e61
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Dec 1 11:49:49 2021 -0800

    Bluetooth: MGMT: Use hci_dev_test_and_{set,clear}_flag

    [ Upstream commit 6f59f991b4e735323f099ac6490e725ae8c750a5 ]
```

```
    This make use of hci_dev_test_and_{set,clear}_flag instead of doing 2
    operations in a row.

    Fixes: cbbdfa6f33198 ("Bluetooth: Enable controller RPA resolution using Experimental
 feature")
    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 35bd12bd9ec73d519cafec00713e9738382cc5e3
 Author: Seevalamuthu Mariappan <quic_seevalam@quicinc.com>
 Date:   Mon Nov 29 16:15:54 2021 +0530

    ath11k: Fix QMI file type enum value

    [ Upstream commit 18ae1ab04525507ae5528245a6df004cacd0d39a ]

    bdf_type for caldata in QMI_WLANFW_BDF_DOWNLOAD_REQ_V01 is wrongly
    sent as 1. But, expected bdf_type value for caldata and EEPROM is 2 and 3
    respectively. It leads to firmware crash. Fix ath11k_qmi_file_type enum
    values.

    Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.5.0.1-01100-QCAHKSWPL_SILICONZ-1
    Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.4.0.1-00192-QCAHKSWPL_SILICONZ-1

    Fixes: 336e7b53c82f ("ath11k: clean up BDF download functions")
    Signed-off-by: Seevalamuthu Mariappan <quic_seevalam@quicinc.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/1638182754-18408-1-git-send-email-
 quic_seevalam@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit fb050f80d5aed1c8628ea4c725611f34c9809cdf
 Author: Fabio Estevam <festevam@denx.de>
 Date:   Wed Nov 24 10:10:47 2021 -0300

    ath10k: Fix the MTU size on QCA9377 SDIO

    [ Upstream commit 09b8cd69edcf2be04a781e1781e98e52a775c9ad ]

    On an imx6dl-pico-pi board with a QCA9377 SDIO chip, simply trying to
    connect via ssh to another machine causes:

    [   55.824159] ath10k_sdio mmc1:0001:1: failed to transmit packet, dropping: -12
    [   55.832169] ath10k_sdio mmc1:0001:1: failed to submit frame: -12
    [   55.838529] ath10k_sdio mmc1:0001:1: failed to push frame: -12
    [   55.905863] ath10k_sdio mmc1:0001:1: failed to transmit packet, dropping: -12
    [   55.913650] ath10k_sdio mmc1:0001:1: failed to submit frame: -12
    [   55.919887] ath10k_sdio mmc1:0001:1: failed to push frame: -12

    , leading to an ssh connection failure.

    One user inspected the size of frames on Wireshark and reported
    the followig:

    "I was able to narrow the issue down to the mtu. If I set the mtu for
    the wlan0 device to 1486 instead of 1500, the issue does not happen.

    The size of frames that I see on Wireshark is exactly 1500 after
    setting it to 1486."

    Clearing the HI_ACS_FLAGS_ALT_DATA_CREDIT_SIZE avoids the problem and
    the ssh command works successfully after that.

    Introduce a 'credit_size_workaround' field to ath10k_hw_params for
```

```
    the QCA9377 SDIO, so that the HI_ACS_FLAGS_ALT_DATA_CREDIT_SIZE
    is not set in this case.

    Tested with QCA9377 SDIO with firmware WLAN.TF.1.1.1-00061-QCATFSWPZ-1.

    Fixes: 2f918ea98606 ("ath10k: enable alt data of TX path for sdio")
    Signed-off-by: Fabio Estevam <festevam@denx.de>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20211124131047.713756-1-festevam@denx.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bdacea482b96d16d71392830c109c8d7ef1948e0
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Tue Dec 7 16:02:41 2021 +0200

    mtd: spi-nor: Fix mtd size for s3an flashes

    [ Upstream commit f656b419d41aabafb6b526abc3988dfbf2e5c1ba ]

    As it was before the blamed commit, s3an_nor_scan() was called
    after mtd size was set with params->size, and it overwrote the mtd
    size value with '8 * nor->page_size * nor->info->n_sectors' when
    XSR_PAGESIZE was set. With the introduction of
    s3an_post_sfdp_fixups(), we missed to update the mtd size for the
    s3an flashes. Fix the mtd size by updating both nor->params->size,
    (which will update the mtd_info size later on) and nor->mtd.size
    (which is used in spi_nor_set_addr_width()).

    Fixes: 641edddb4f43 ("mtd: spi-nor: Add s3an_post_sfdp_fixups()")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Link: https://lore.kernel.org/r/20211207140254.87681-2-tudor.ambarus@microchip.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29a0950b36bd6892a82bfa123bfadeb3fe60d160
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Fri Oct 29 20:26:12 2021 +0300

    mtd: spi-nor: Get rid of nor->page_size

    [ Upstream commit 5854d4a6cc356ba3e16d8593ac1c089a32d1759c ]

    nor->page_size duplicated what nor->params->page_size indicates
    for no good reason. page_size is a flash parameter of fixed value
    and it is better suited to be found in nor->params->page_size.

    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Reviewed-by: Michael Walle <michael@walle.cc>
    Link: https://lore.kernel.org/r/20211029172633.886453-5-tudor.ambarus@microchip.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4b5d9e98ba0d3247f9ffb482c0c59162b82a03bc
Author: Li Hua <hucool.lihua@huawei.com>
Date:   Fri Dec 3 03:36:18 2021 +0000

    sched/rt: Try to restart rt period timer when rt runtime exceeded

    [ Upstream commit 9b58e976b3b391c0cf02e038d53dd0478ed3013c ]

    When rt_runtime is modified from -1 to a valid control value, it may
    cause the task to be throttled all the time. Operations like the following
    will trigger the bug. E.g:

      1. echo -1 > /proc/sys/kernel/sched_rt_runtime_us
```

```
       2. Run a FIFO task named A that executes while(1)
       3. echo 950000 > /proc/sys/kernel/sched_rt_runtime_us

     When rt_runtime is -1, The rt period timer will not be activated when task
     A enqueued. And then the task will be throttled after setting rt_runtime to
     950,000. The task will always be throttled because the rt period timer is
     not activated.

     Fixes: d0b27fa77854 ("sched: rt-group: synchonised bandwidth period")
     Reported-by: Hulk Robot <hulkci@huawei.com>
     Signed-off-by: Li Hua <hucool.lihua@huawei.com>
     Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
     Link: https://lkml.kernel.org/r/20211203033618.11895-1-hucool.lihua@huawei.com
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b570c321d4062e141dad4240522fbf17270a22c0
Author: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
Date:   Fri Apr 2 22:47:55 2021 -0700

    wireless: iwlwifi: Fix a double free in iwl_txq_dyn_alloc_dma

    [ Upstream commit f973795a8d19cbf3d03807704eb7c6ff65788d5a ]

    In iwl_txq_dyn_alloc_dma, txq->tfds is freed at first time by:
    iwl_txq_alloc()->goto err_free_tfds->dma_free_coherent(). But
    it forgot to set txq->tfds to NULL.

    Then the txq->tfds is freed again in iwl_txq_dyn_alloc_dma by:
    goto error->iwl_txq_gen2_free_memory()->dma_free_coherent().

    My patch sets txq->tfds to NULL after the first free to avoid the
    double free.

    Fixes: 0cd1ad2d7fd41 ("iwlwifi: move all bus-independent TX functions to common code")
    Signed-off-by: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
    Link: https://lore.kernel.org/r/20210403054755.4781-1-lyl2019@mail.ustc.edu.cn
    Signed-off-by: Luca Coelho <luciano.coelho@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 818ae17fe90a92b6b4815db150ca45fe8040efc4
Author: Robert Schlabbach <robert_s@gmx.net>
Date:   Wed Dec 1 22:08:43 2021 +0100

    media: si2157: Fix "warm" tuner state detection

    [ Upstream commit a6441ea29cb2c9314654e093a1cd8020b9b851c8 ]

    Commit e955f959ac52 ("media: si2157: Better check for running tuner in
    init") completely broke the "warm" tuner detection of the si2157 driver
    due to a simple endian error: The Si2157 CRYSTAL_TRIM property code is
    0x0402 and needs to be transmitted LSB first. However, it was inserted
    MSB first, causing the warm detection to always fail and spam the kernel
    log with tuner initialization messages each time the DVB frontend
    device was closed and reopened:

    [  312.215682] si2157 16-0060: found a 'Silicon Labs Si2157-A30'
    [  312.264334] si2157 16-0060: firmware version: 3.0.5
    [  342.248593] si2157 16-0060: found a 'Silicon Labs Si2157-A30'
    [  342.295743] si2157 16-0060: firmware version: 3.0.5
    [  372.328574] si2157 16-0060: found a 'Silicon Labs Si2157-A30'
    [  372.385035] si2157 16-0060: firmware version: 3.0.5

    Also, the reinitializations were observed disturb _other_ tuners on
    multi-tuner cards such as the Hauppauge WinTV-QuadHD, leading to missed
    or errored packets when one of the other DVB frontend devices on that
```

```
    card was opened.

    Fix the order of the property code bytes to make the warm detection work
    again, also reducing the tuner initialization message in the kernel log
    to once per power-on, as well as fixing the interference with other
    tuners.

    Link: https://lore.kernel.org/linux-media/trinity-2a86eb9d-6264-4387-95e1-ba7b79a4050f-
1638392923493@3c-app-gmx-bap03

    Fixes: e955f959ac52 ("media: si2157: Better check for running tuner in init")
    Reported-by: Robert Schlabbach <robert_s@gmx.net>
    Signed-off-by: Robert Schlabbach <robert_s@gmx.net>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35a3dfe1a167bbc0bca869bf7cd3c02558ff3aa3
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Tue Nov 30 17:34:44 2021 +0100

    media: saa7146: mxb: Fix a NULL pointer dereference in mxb_attach()

    [ Upstream commit 0407c49ebe330333478440157c640fffd986f41b ]

    In mxb_attach(dev, info), saa7146_vv_init() is called to allocate a
    new memory for dev->vv_data. saa7146_vv_release() will be called on
    failure of mxb_probe(dev). There is a dereference of dev->vv_data
    in saa7146_vv_release(), which could lead to a NULL pointer dereference
    on failure of saa7146_vv_init().

    Fix this bug by adding a check of saa7146_vv_init().

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_VIDEO_MXB=m show no new warnings,
    and our static analyzer no longer warns about this code.

    Fixes: 03b1930efd3c ("V4L/DVB: saa7146: fix regression of the av7110/budget-av driver")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f486f70b78768d3d25ee65a6d5e193cab7ad6a68
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Tue Nov 30 16:38:05 2021 +0100

    media: dib8000: Fix a memleak in dib8000_init()

    [ Upstream commit 8dbdcc7269a83305ee9d677b75064d3530a48ee2 ]

    In dib8000_init(), the variable fe is not freed or passed out on the
    failure of dib8000_identify(&state->i2c), which could lead to a memleak.

    Fix this bug by adding a kfree of fe in the error path.

    This bug was found by a static analyzer. The analysis employs
```

```
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_DVB_DIB8000=m show no new warnings,
    and our static analyzer no longer warns about this code.

    Fixes: 77e2c0f5d471 ("V4L/DVB (12900): DiB8000: added support for DiBcom ISDB-T/ISDB-Tsb
demodulator DiB8000")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48c61680c1c8a65d9fee944f965d589d04357fc9
Author: Ricardo Ribalda <ribalda@chromium.org>
Date:   Wed Dec 1 06:22:18 2021 +0100

    media: uvcvideo: Avoid returning invalid controls

    [ Upstream commit 414d3b49d9fd4a0bb16a13d929027847fd094f3f ]

    If the memory where ctrl_found is placed has the value of uvc_ctrl and
    __uvc_find_control does not find the control we will return an invalid
    index.

    Fixes: 6350d6a4ed487 ("media: uvcvideo: Set error_idx during ctrl_commit errors")
    Signed-off-by: Ricardo Ribalda <ribalda@chromium.org>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e40f71587469543a68986c6f26c6e4ba1e3d02e5
Author: Ricardo Ribalda <ribalda@chromium.org>
Date:   Wed Dec 1 06:22:17 2021 +0100

    media: uvcvideo: Avoid invalid memory access

    [ Upstream commit f0577b1b6394f954903fcc67e12fe9e7001dafd6 ]

    If mappings points to an invalid memory, we will be invalid accessing
    it. Solve it by initializing the value of the variable mapping and by
    changing the order in the conditional statement (to avoid accessing
    mapping->id if not needed).

    Fix:
    kasan: GPF could be caused by NULL-ptr deref or user memory access
    general protection fault: 0000 [#1] PREEMPT SMP KASAN NOPTI

    Fixes: 6350d6a4ed487 ("media: uvcvideo: Set error_idx during ctrl_commit errors")
    Signed-off-by: Ricardo Ribalda <ribalda@chromium.org>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5ece5b66f554e824e583f7ec160d607334c64a83
Author: Colin Ian King <colin.king@intel.com>
Date:   Fri Sep 17 13:49:30 2021 +0200

    media: uvcvideo: Fix memory leak of object map on error exit path
```

```
    [ Upstream commit 4b065060555b14c7a9b86c013a1c9bee8e8b6fbd ]

    Currently when the allocation of map->name fails the error exit path
    does not kfree the previously allocated object map. Fix this by
    setting ret to -ENOMEM and taking the free_map exit error path to
    ensure map is kfree'd.

    Addresses-Coverity: ("Resource leak")

    Fixes: 70fa906d6fce ("media: uvcvideo: Use control names from framework")
    Signed-off-by: Colin Ian King <colin.king@canonical.com>
    Reviewed-by: Ricardo Ribalda <ribalda@chromium.org>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eabdce1bf11801a1e159ead65e20d8c6c9c28901
Author: Alexander Lobakin <alexandr.lobakin@intel.com>
Date:   Fri Dec 3 20:50:04 2021 +0100

    samples: bpf: Fix 'unknown warning group' build warning on Clang

    [ Upstream commit 6f670d06e47c774bc065aaa84a527a4838f34bd8 ]

    Clang doesn't have 'stringop-truncation' group like GCC does, and
    complains about it when building samples which use xdp_sample_user
    infra:

     samples/bpf/xdp_sample_user.h:48:32: warning: unknown warning group '-Wstringop-
truncation', ignored [-Wunknown-warning-option]
     #pragma GCC diagnostic ignored "-Wstringop-truncation"
                                    ^
    [ repeat ]

    Those are harmless, but avoidable when guarding it with ifdef.
    I could guard push/pop as well, but this would require one more
    ifdef cruft around a single line which I don't think is reasonable.

    Fixes: 156f886cf697 ("samples: bpf: Add basic infrastructure for XDP samples")
    Signed-off-by: Alexander Lobakin <alexandr.lobakin@intel.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/bpf/20211203195004.5803-3-alexandr.lobakin@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6e6f5b3ff4189da2e506137ee7fd227fcfdaf439
Author: Alexander Lobakin <alexandr.lobakin@intel.com>
Date:   Fri Dec 3 20:50:03 2021 +0100

    samples: bpf: Fix xdp_sample_user.o linking with Clang

    [ Upstream commit e64fbcaa7a666f16329b1c67af15ea501bc84586 ]

    Clang (13) doesn't get the jokes about specifying libraries to link in
    cclags of individual .o objects:

    clang-13: warning: -lm: 'linker' input unused [-Wunused-command-line-argument]
    [ ... ]
      LD  samples/bpf/xdp_redirect_cpu
      LD  samples/bpf/xdp_redirect_map_multi
      LD  samples/bpf/xdp_redirect_map
      LD  samples/bpf/xdp_redirect
      LD  samples/bpf/xdp_monitor
    /usr/bin/ld: samples/bpf/xdp_sample_user.o: in function `sample_summary_print':
```

```
    xdp_sample_user.c:(.text+0x84c): undefined reference to `floor'
    /usr/bin/ld: xdp_sample_user.c:(.text+0x870): undefined reference to `ceil'
    /usr/bin/ld: xdp_sample_user.c:(.text+0x8cf): undefined reference to `floor'
    /usr/bin/ld: xdp_sample_user.c:(.text+0x8f3): undefined reference to `ceil'
    [ more ]

    Specify '-lm' as ldflags for all xdp_sample_user.o users in the main
    Makefile and remove it from ccflags of ^ in Makefile.target -- just
    like it's done for all other samples. This works with all compilers.

    Fixes: 6e1051a54e31 ("samples: bpf: Convert xdp_monitor to XDP samples helper")
    Fixes: b926c55d856c ("samples: bpf: Convert xdp_redirect to XDP samples helper")
    Fixes: e531a220cc59 ("samples: bpf: Convert xdp_redirect_cpu to XDP samples helper")
    Fixes: bbe65865aa05 ("samples: bpf: Convert xdp_redirect_map to XDP samples helper")
    Fixes: 594a116b2aa1 ("samples: bpf: Convert xdp_redirect_map_multi to XDP samples
helper")
    Signed-off-by: Alexander Lobakin <alexandr.lobakin@intel.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/bpf/20211203195004.5803-2-alexandr.lobakin@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1912d4470abe48393ae7c95be0bd767d1da316f1
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Dec 1 15:28:22 2021 -0800

    samples/bpf: Clean up samples/bpf build failes

    [ Upstream commit 527024f7aeb683ce7ef49b07ef7ce9ecf015288d ]

    Remove xdp_samples_user.o rule redefinition which generates Makefile
    warning and instead override TPROGS_CFLAGS. This seems to work fine when
    building inside selftests/bpf.

    That was one big head-scratcher before I found that generic
    Makefile.target hid this surprising specialization for for xdp_samples_user.o.

    Main change is to use actual locally installed libbpf headers.

    Also drop printk macro re-definition (not even used!).

    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Link: https://lore.kernel.org/bpf/20211201232824.3166325-8-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f6d3ccce02a428aff98a71974052f3a087367462
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Nov 28 20:00:28 2021 +0100

    platform/x86: wmi: Fix driver->notify() vs ->probe() race

    [ Upstream commit 9918878676a5f9e99b98679f04b9e6c0f5426b0a ]

    The driver core sets struct device->driver before calling out
    to the bus' probe() method, this leaves a window where an ACPI
    notify may happen on the WMI object before the driver's
    probe() method has completed running, causing e.g. the
    driver's notify() callback to get called with drvdata
    not yet being set leading to a NULL pointer deref.

    At a check for this to the WMI core, ensuring that the notify()
    callback is not called before the driver is ready.

    Fixes: 1686f5444546 ("platform/x86: wmi: Incorporate acpi_install_notify_handler")
```

```
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link: https://lore.kernel.org/r/20211128190031.405620-2-hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 36673b9d036e42677d3a29833c388ece43cb00a4
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Nov 28 20:00:27 2021 +0100

    platform/x86: wmi: Replace read_takes_no_args with a flags field

    [ Upstream commit a90b38c58667142ecff2521481ed44286d46b140 ]

    Replace the wmi_block.read_takes_no_args bool field with
    an unsigned long flags field, used together with test_bit()
    and friends.

    This is a preparation patch for fixing a driver->notify() vs ->probe()
    race, which requires atomic flag handling.

    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link: https://lore.kernel.org/r/20211128190031.405620-1-hdegoede@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e29bfcac41205f3b43f9e7e38d2e66d6143d7916
Author: Reiji Watanabe <reijiw@google.com>
Date:   Sun Dec 5 16:47:36 2021 -0800

    arm64: mte: DC {GVA,GZVA} shouldn't be used when DCZID_EL0.DZP == 1

    [ Upstream commit 685e2564daa1493053fcd7f1dbed38b35ee2f3cb ]

    Currently, mte_set_mem_tag_range() and mte_zero_clear_page_tags() use
    DC {GVA,GZVA} unconditionally.  But, they should make sure that
    DCZID_EL0.DZP, which indicates whether or not use of those instructions
    is prohibited, is zero when using those instructions.
    Use ST{G,ZG,Z2G} instead when DCZID_EL0.DZP == 1.

    Fixes: 013bb59dbb7c ("arm64: mte: handle tags zeroing at page allocation time")
    Fixes: 3d0cca0b02ac ("kasan: speed up mte_set_mem_tag_range")
    Signed-off-by: Reiji Watanabe <reijiw@google.com>
    Link: https://lore.kernel.org/r/20211206004736.1520989-3-reijiw@google.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06ff0e34ecc142dd0a050092238bb87523417244
Author: Reiji Watanabe <reijiw@google.com>
Date:   Sun Dec 5 16:47:35 2021 -0800

    arm64: clear_page() shouldn't use DC ZVA when DCZID_EL0.DZP == 1

    [ Upstream commit f0616abd4e67143b45b04b565839148458857347 ]

    Currently, clear_page() uses DC ZVA instruction unconditionally.  But it
    should make sure that DCZID_EL0.DZP, which indicates whether or not use
    of DC ZVA instruction is prohibited, is zero when using the instruction.
    Use STNP instead when DCZID_EL0.DZP == 1.

    Fixes: f27bb139c387 ("arm64: Miscellaneous library functions")
    Signed-off-by: Reiji Watanabe <reijiw@google.com>
    Reviewed-by: Robin Murphy <robin.murphy@arm.com>
    Link: https://lore.kernel.org/r/20211206004736.1520989-2-reijiw@google.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 1f10491e88993957900f5f35978edea65378b28d
Author: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
Date:   Fri Nov 26 09:35:14 2021 +0000

    drm: rcar-du: crtc: Support external DSI dot clock

    [ Upstream commit 57b290cb905bec520372ac635d9e9f0548d9d67e ]

    On platforms with an external clock, both the group and crtc must be
    handled accordingly to correctly pass through the external clock and
    configure the DU to use the external rate.

    The CRTC support was missed while adding the DSI support on the r8a779a0
    which led to the output clocks being incorrectly determined.

    Ensure that when a CRTC is routed through the DSI encoder, the external
    clock is used without any further divider being applied.

    Fixes: b291fdcf5114 ("drm: rcar-du: Add r8a779a0 device support")
    Signed-off-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3f31e92e6bbedafe15e87b5b281b7b7ec229eada
Author: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
Date:   Mon Nov 29 17:08:45 2021 +0000

    drm: rcar-du: Add DSI support to rcar_du_output_name

    [ Upstream commit e0e4c64a64780a8672480618142776de8bf98d07 ]

    The DSI output names were not added when the DSI pipeline support was
    introduced.

    Add the correct labels for these outputs, and fix the sort order to
    match 'enum rcar_du_output' while we are here.

    Fixes: b291fdcf5114 ("drm: rcar-du: Add r8a779a0 device support")
    Suggested-by: Biju Das <biju.das.jz@bp.renesas.com>
    Signed-off-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fd8097e7adc91756b50158e4de2345771505165
Author: Kajol Jain <kjain@linux.ibm.com>
Date:   Mon Dec 6 13:03:15 2021 +0530

    bpf: Remove config check to enable bpf support for branch records

    [ Upstream commit db52f57211b4e45f0ebb274e2c877b211dc18591 ]

    Branch data available to BPF programs can be very useful to get stack traces
    out of userspace application.

    Commit fff7b64355ea ("bpf: Add bpf_read_branch_records() helper") added BPF
    support to capture branch records in x86. Enable this feature also for other
    architectures as well by removing checks specific to x86.

    If an architecture doesn't support branch records, bpf_read_branch_records()
    still has appropriate checks and it will return an -EINVAL in that scenario.
    Based on UAPI helper doc in include/uapi/linux/bpf.h, unsupported architectures
    should return -ENOENT in such case. Hence, update the appropriate check to
```

```
    return -ENOENT instead.

    Selftest 'perf_branches' result on power9 machine which has the branch stacks
    support:

     - Before this patch:

      [command]# ./test_progs -t perf_branches
       #88/1 perf_branches/perf_branches_hw:FAIL
       #88/2 perf_branches/perf_branches_no_hw:OK
       #88 perf_branches:FAIL
      Summary: 0/1 PASSED, 0 SKIPPED, 1 FAILED

     - After this patch:

      [command]# ./test_progs -t perf_branches
       #88/1 perf_branches/perf_branches_hw:OK
       #88/2 perf_branches/perf_branches_no_hw:OK
       #88 perf_branches:OK
      Summary: 1/2 PASSED, 0 SKIPPED, 0 FAILED

    Selftest 'perf_branches' result on power9 machine which doesn't have branch
    stack report:

     - After this patch:

      [command]# ./test_progs -t perf_branches
       #88/1 perf_branches/perf_branches_hw:SKIP
       #88/2 perf_branches/perf_branches_no_hw:OK
       #88 perf_branches:OK
      Summary: 1/1 PASSED, 1 SKIPPED, 0 FAILED

    Fixes: fff7b64355eac ("bpf: Add bpf_read_branch_records() helper")
    Suggested-by: Peter Zijlstra <peterz@infradead.org>
    Signed-off-by: Kajol Jain <kjain@linux.ibm.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211206073315.77432-1-kjain@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 80901137a50a7924390b3afe52460ca54c0b5f33
Author: Hou Tao <houtao1@huawei.com>
Date:   Fri Dec 3 13:30:01 2021 +0800

    bpf: Disallow BPF_LOG_KERNEL log level for bpf(BPF_BTF_LOAD)

    [ Upstream commit 866de407444398bc8140ea70de1dba5f91cc34ac ]

    BPF_LOG_KERNEL is only used internally, so disallow bpf_btf_load()
    to set log level as BPF_LOG_KERNEL. The same checking has already
    been done in bpf_check(), so factor out a helper to check the
    validity of log attributes and use it in both places.

    Fixes: 8580ac9404f6 ("bpf: Process in-kernel BTF")
    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Acked-by: Martin KaFai Lau <kafai@fb.com>
    Link: https://lore.kernel.org/bpf/20211203053001.740945-1-houtao1@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00a07270bf39b560cbf0f66ba4ccda944f389720
Author: Alexei Starovoitov <ast@kernel.org>
Date:   Wed Dec 1 10:10:29 2021 -0800

    bpf: Adjust BTF log size limit.
```

```
    [ Upstream commit c5a2d43e998a821701029f23e25b62f9188e93ff ]

    Make BTF log size limit to be the same as the verifier log size limit.
    Otherwise tools that progressively increase log size and use the same log
    for BTF loading and program loading will be hitting hard to debug EINVAL.

    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20211201181040.23337-7-alexei.starovoitov@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b481d22656754b8d83141123d9aabc9739ced852
Author: Vincent Donnefort <vincent.donnefort@arm.com>
Date:   Mon Nov 29 17:31:15 2021 +0000

    sched/fair: Fix per-CPU kthread and wakee stacking for asym CPU capacity

    [ Upstream commit 014ba44e8184e1acf93e0cbb7089ee847802f8f0 ]

    select_idle_sibling() has a special case for tasks woken up by a per-CPU
    kthread where the selected CPU is the previous one. For asymmetric CPU
    capacity systems, the assumption was that the wakee couldn't have a
    bigger utilization during task placement than it used to have during the
    last activation. That was not considering uclamp.min which can completely
    change between two task activations and as a consequence mandates the
    fitness criterion asym_fits_capacity(), even for the exit path described
    above.

    Fixes: b4c9c9f15649 ("sched/fair: Prefer prev cpu in asymmetric wakeup path")
    Signed-off-by: Vincent Donnefort <vincent.donnefort@arm.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Valentin Schneider <valentin.schneider@arm.com>
    Reviewed-by: Dietmar Eggemann <dietmar.eggemann@arm.com>
    Link: https://lkml.kernel.org/r/20211129173115.4006346-1-vincent.donnefort@arm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb77072cf2e739b9259e7af9561d76f8ca967e0f
Author: Vincent Donnefort <vincent.donnefort@arm.com>
Date:   Wed Dec 1 14:34:50 2021 +0000

    sched/fair: Fix detection of per-CPU kthreads waking a task

    [ Upstream commit 8b4e74ccb582797f6f0b0a50372ebd9fd2372a27 ]

    select_idle_sibling() has a special case for tasks woken up by a per-CPU
    kthread, where the selected CPU is the previous one. However, the current
    condition for this exit path is incomplete. A task can wake up from an
    interrupt context (e.g. hrtimer), while a per-CPU kthread is running. A
    such scenario would spuriously trigger the special case described above.
    Also, a recent change made the idle task like a regular per-CPU kthread,
    hence making that situation more likely to happen
    (is_per_cpu_kthread(swapper) being true now).

    Checking for task context makes sure select_idle_sibling() will not
    interpret a wake up from any other context as a wake up by a per-CPU
    kthread.

    Fixes: 52262ee567ad ("sched/fair: Allow a per-CPU kthread waking a task to stack on the
same CPU, to fix XFS performance regression")
    Signed-off-by: Vincent Donnefort <vincent.donnefort@arm.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Vincent Guittot <vincent.guittot@linaro.org>
    Reviewed-by: Valentin Schneider <valentin.schneider@arm.com>
    Link: https://lore.kernel.org/r/20211201143450.479472-1-vincent.donnefort@arm.com
```

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

commit efdbc07e0da4475d061383fc4858065d4aeb866a
Author: Maxim Mikityanskiy <maximmi@nvidia.com>
Date:   Tue Nov 30 20:18:11 2021 +0200

    bpf: Fix the test_task_vma selftest to support output shorter than 1 kB

    [ Upstream commit da54ab14953c38d98cb3e34c564c06c3739394b2 ]

    The test for bpf_iter_task_vma assumes that the output will be longer
    than 1 kB, as the comment above the loop says. Due to this assumption,
    the loop becomes infinite if the output turns to be shorter than 1 kB.
    The return value of read_fd_into_buffer is 0 when the end of file was
    reached, and len isn't being increased any more.

    This commit adds a break on EOF to handle short output correctly. For
    the reference, this is the contents that I get when running test_progs
    under vmtest.sh, and it's shorter than 1 kB:

    00400000-00401000 r--p 00000000 fe:00 25867     /root/bpf/test_progs
    00401000-00674000 r-xp 00001000 fe:00 25867     /root/bpf/test_progs
    00674000-0095f000 r--p 00274000 fe:00 25867     /root/bpf/test_progs
    0095f000-00983000 r--p 0055e000 fe:00 25867     /root/bpf/test_progs
    00983000-00a8a000 rw-p 00582000 fe:00 25867     /root/bpf/test_progs
    00a8a000-0484e000 rw-p 00000000 00:00 0
    7f6c64000000-7f6c64021000 rw-p 00000000 00:00 0
    7f6c64021000-7f6c68000000 ---p 00000000 00:00 0
    7f6c6ac8f000-7f6c6ac90000 r--s 00000000 00:0d 8032
    anon_inode:bpf-map
    7f6c6ac90000-7f6c6ac91000 ---p 00000000 00:00 0
    7f6c6ac91000-7f6c6b491000 rw-p 00000000 00:00 0
    7f6c6b491000-7f6c6b492000 r--s 00000000 00:0d 8032
    anon_inode:bpf-map
    7f6c6b492000-7f6c6b493000 rw-s 00000000 00:0d 8032
    anon_inode:bpf-map
    7ffc1e23d000-7ffc1e25e000 rw-p 00000000 00:00 0
    7ffc1e3b8000-7ffc1e3bc000 r--p 00000000 00:00 0
    7ffc1e3bc000-7ffc1e3bd000 r-xp 00000000 00:00 0
    7fffffffe000-7ffffffff000 --xp 00000000 00:00 0

    Fixes: e8168840e16c ("selftests/bpf: Add test for bpf_iter_task_vma")
    Signed-off-by: Maxim Mikityanskiy <maximmi@nvidia.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211130181811.594220-1-maximmi@nvidia.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9b792ab3705df05a0bc3f661bd42eee6d3c80e9
Author: Sean Wang <sean.wang@mediatek.com>
Date:   Thu Dec 2 02:02:47 2021 +0800

    Bluetooth: btmtksdio: fix resume failure

    [ Upstream commit 561ae1d46a8ddcbc13162d5771f5ed6c8249e730 ]

    btmtksdio have to rely on MMC_PM_KEEP_POWER in pm_flags to avoid that
    SDIO power is being shut off during the device is in suspend. That fixes
    the SDIO command fails to access the bus after the device is resumed.

    Fixes: 7f3c563c575e7 ("Bluetooth: btmtksdio: Add runtime PM support to SDIO based
Bluetooth")
    Co-developed-by: Mark-yw Chen <mark-yw.chen@mediatek.com>
    Signed-off-by: Mark-yw Chen <mark-yw.chen@mediatek.com>
    Signed-off-by: Sean Wang <sean.wang@mediatek.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
```

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

commit 103ad61f47de51875ace0d3ced45b45d7ce37b55
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Dec 2 11:07:03 2021 +0800

    staging: rtl8192e: rtllib_module: fix error handle case in alloc_rtllib()

    [ Upstream commit e730cd57ac2dfe94bca0f14a3be8e1b21de41a9c ]

    Some variables are leaked in the error handling in alloc_rtllib(), free
    the variables in the error path.

    Fixes: 94a799425eee ("From: wlanfae <wlanfae@realtek.com>")
    Reviewed-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20211202030704.2425621-3-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1843aee812f0b7d4afee31d61412cc4ca16f9e88
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Dec 2 11:07:02 2021 +0800

    staging: rtl8192e: return error code from rtllib_softmac_init()

    [ Upstream commit 68bf78ff59a0891eb1239948e94ce10f73a9dd30 ]

    If it fails to allocate 'dot11d_info', rtllib_softmac_init()
    should return error code. And remove unneccessary error message.

    Fixes: 94a799425eee ("From: wlanfae <wlanfae@realtek.com>")
    Reviewed-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20211202030704.2425621-2-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6a8b9d4f13809475f04184caa4532a733b0f0c1e
Author: Tasos Sahanidis <tasos@tasossah.com>
Date:   Fri Sep 3 09:47:58 2021 +0300

    floppy: Fix hang in watchdog when disk is ejected

    [ Upstream commit fb48febce7e30baed94dd791e19521abd2c3fd83 ]

    When the watchdog detects a disk change, it calls cancel_activity(),
    which in turn tries to cancel the fd_timer delayed work.

    In the above scenario, fd_timer_fn is set to fd_watchdog(), meaning
    it is trying to cancel its own work.
    This results in a hang as cancel_delayed_work_sync() is waiting for the
    watchdog (itself) to return, which never happens.

    This can be reproduced relatively consistently by attempting to read a
    broken floppy, and ejecting it while IO is being attempted and retried.

    To resolve this, this patch calls cancel_delayed_work() instead, which
    cancels the work without waiting for the watchdog to return and finish.

    Before this regression was introduced, the code in this section used
    del_timer(), and not del_timer_sync() to delete the watchdog timer.
```

```
    Link: https://lore.kernel.org/r/399e486c-6540-db27-76aa-7a271b061f76@tasossah.com
    Fixes: 070ad7e793dc ("floppy: convert to delayed work and single-thread wq")
    Signed-off-by: Tasos Sahanidis <tasos@tasossah.com>
    Signed-off-by: Denis Efremov <efremov@linux.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c406ed2650394a9bf61beb00c6a36f19c3deaeff
Author: Michael Walle <michael@walle.cc>
Date:   Thu Nov 4 14:48:43 2021 +0100

    mtd: core: provide unique name for nvmem device

    [ Upstream commit c048b60d39e109c201d31ed5ad3a4f939064d6c4 ]

    If there is more than one mtd device which supports OTP, there will
    be a kernel warning about duplicated sysfs entries and the probing will
    fail. This is because the nvmem device name is not unique. Make it
    unique by prepending the name of the mtd. E.g. before the name was
    "user-otp", now it will be "mtd0-user-otp".

    For reference the kernel splash is:
    [    4.665531] sysfs: cannot create duplicate filename '/bus/nvmem/devices/user-otp'
    [    4.673056] CPU: 0 PID: 1 Comm: swapper/0 Not tainted 5.15.0-next-20211101+ #1296
    [    4.680565] Hardware name: Kontron SMARC-sAL28 (Single PHY) on SMARC Eval 2.0 carrier
(DT)
    [    4.688856] Call trace:
    [    4.691303]  dump_backtrace+0x0/0x1bc
    [    4.694984]  show_stack+0x24/0x30
    [    4.698306]  dump_stack_lvl+0x68/0x84
    [    4.701980]  dump_stack+0x18/0x34
    [    4.705302]  sysfs_warn_dup+0x70/0x90
    [    4.708973]  sysfs_do_create_link_sd+0x144/0x150
    [    4.713603]  sysfs_create_link+0x2c/0x50
    [    4.717535]  bus_add_device+0x74/0x120
    [    4.721293]  device_add+0x330/0x890
    [    4.724791]  device_register+0x2c/0x40
    [    4.728550]  nvmem_register+0x240/0x9f0
    [    4.732398]  mtd_otp_nvmem_register+0xb0/0x10c
    [    4.736854]  mtd_device_parse_register+0x28c/0x2b4
    [    4.741659]  spi_nor_probe+0x20c/0x2e0
    [    4.745418]  spi_mem_probe+0x78/0xbc
    [    4.749001]  spi_probe+0x90/0xf0
    [    4.752237]  really_probe.part.0+0xa4/0x320
    ..
    [    4.873936] mtd mtd1: Failed to register OTP NVMEM device
    [    4.894468] spi-nor: probe of spi0.0 failed with error -17

    Fixes: 4b361cfa8624 ("mtd: core: add OTP nvmem provider support")
    Signed-off-by: Michael Walle <michael@walle.cc>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211104134843.2642800-1-michael@walle.cc
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a688d983cc03f469609ba9e4c67f5b5cd28d6d6
Author: Lino Sanfilippo <LinoSanfilippo@gmx.de>
Date:   Mon Nov 29 18:42:38 2021 +0100

    serial: amba-pl011: do not request memory region twice

    [ Upstream commit d1180405c7b5c7a1c6bde79d5fc24fe931430737 ]

    With commit 3873e2d7f63a ("drivers: PL011: refactor pl011_probe()") the
    function devm_ioremap() called from pl011_setup_port() was replaced with
    devm_ioremap_resource(). Since this function not only remaps but also
```

```
    requests the ports io memory region it now collides with the .config_port()
    callback which requests the same region at uart port registration.

    Since devm_ioremap_resource() already claims the memory successfully, the
    request in .config_port() fails.

    Later at uart port deregistration the attempt to release the unclaimed
    memory also fails. The failure results in a “Trying to free nonexistent
    resource" warning.

    Fix these issues by removing the callbacks that implement the redundant
    memory allocation/release. Also make sure that changing the drivers io
    memory base address via TIOCSSERIAL is not allowed any more.

    Fixes: 3873e2d7f63a ("drivers: PL011: refactor pl011_probe()")
    Signed-off-by: Lino Sanfilippo <LinoSanfilippo@gmx.de>
    Link: https://lore.kernel.org/r/20211129174238.8333-1-LinoSanfilippo@gmx.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 22215d99618ac1de095d0068f6f97369105e9769
Author: Lizhi Hou <lizhi.hou@xilinx.com>
Date:   Mon Nov 29 12:23:02 2021 -0800

    tty: serial: uartlite: allow 64 bit address

    [ Upstream commit 3672fb65155530b5eea6225685c75329b6debec3 ]

    The base address of uartlite registers could be 64 bit address which is from
    device resource. When ulite_probe() calls ulite_assign(), this 64 bit
    address is casted to 32-bit. The fix is to replace "u32" type with
    "phys_addr_t" type for the base address in ulite_assign() argument list.

    Fixes: 8fa7b6100693 ("[POWERPC] Uartlite: Separate the bus binding from the driver
proper")
    Signed-off-by: Lizhi Hou <lizhi.hou@xilinx.com>
    Link: https://lore.kernel.org/r/20211129202302.1319033-1-lizhi.hou@xilinx.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b2ac3c06305e4b1bf8d74faee570844036735950
Author: Nishanth Menon <nm@ti.com>
Date:   Fri Nov 12 22:26:40 2021 -0600

    arm64: dts: ti: k3-j7200: Correct the d-cache-sets info

    [ Upstream commit a172c86931709d6663318609d71a811333bdf4b0 ]

    A72 Cluster (chapter 1.3.1 [1]) has 48KB Icache, 32KB Dcache and 1MB L2 Cache
     - ICache is 3-way set-associative
     - Dcache is 2-way set-associative
     - Line size are 64bytes

    32KB (Dcache)/64 (fixed line length of 64 bytes) = 512 ways
    512 ways / 2 (Dcache is 2-way per set) = 256 sets.

    So, correct the d-cache-sets info.

    [1] https://www.ti.com/lit/pdf/spruiu1

    Fixes: d361ed88455f ("arm64: dts: ti: Add support for J7200 SoC")
    Reported-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Reviewed-by: Kishon Vijay Abraham I <kishon@ti.com>
```

```
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211113042640.30955-1-nm@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ec0614cff1488460a6d1a32f23e2b88108b57cd
Author: Nishanth Menon <nm@ti.com>
Date:   Fri Nov 12 22:36:39 2021 -0600

    arm64: dts: ti: k3-j721e: Fix the L2 cache sets

    [ Upstream commit e9ba3a5bc6fdc2c796c69fdaf5ed6c9957cf9f9d ]

    A72's L2 cache[1] on J721e[2] is 1MB. A72's L2 is fixed line length of
    64 bytes and 16-way set-associative cache structure.

    1MB of L2 / 64 (line length) = 16384 ways
    16384 ways / 16 = 1024 sets

    Fix the l2 cache-sets.

    [1] https://developer.arm.com/documentation/100095/0003/Level-2-Memory-System/About-the-
L2-memory-system
    [2] http://www.ti.com/lit/pdf/spruil1

    Fixes: 2d87061e70de ("arm64: dts: ti: Add Support for J721E SoC")
    Reported-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211113043639.4413-1-nm@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b4313b1f4b1d9524effd9d3720ab06f30840a51
Author: Nishanth Menon <nm@ti.com>
Date:   Fri Nov 12 22:36:38 2021 -0600

    arm64: dts: ti: k3-j7200: Fix the L2 cache sets

    [ Upstream commit d0c826106f3fc11ff97285102b576b65576654ae ]

    A72's L2 cache[1] on J7200[2] is 1MB. A72's L2 is fixed line length of
    64 bytes and 16-way set-associative cache structure.

    1MB of L2 / 64 (line length) = 16384 ways
    16384 ways / 16 = 1024 sets

    Fix the l2 cache-sets.

    [1] https://developer.arm.com/documentation/100095/0003/Level-2-Memory-System/About-the-
L2-memory-system
    [2] https://www.ti.com/lit/pdf/spruiu1

    Fixes: d361ed88455f ("arm64: dts: ti: Add support for J7200 SoC")
    Reported-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211113043638.4358-1-nm@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03eee947a6268eaa031ec4bad6f0de83f3c78ae3
Author: Nishanth Menon <nm@ti.com>
Date:   Fri Nov 12 22:36:35 2021 -0600

    arm64: dts: ti: k3-am642: Fix the L2 cache sets
```

```
    [ Upstream commit a27a93bf70045be54b594fa8482959ffb84166d7 ]

    A53's L2 cache[1] on AM642[2] is 256KB. A53's L2 is fixed line length
    of 64 bytes and 16-way set-associative cache structure.

    256KB of L2 / 64 (line length) = 4096 ways
    4096 ways / 16 = 256 sets

    Fix the l2 cache-sets.

    [1] https://developer.arm.com/documentation/ddi0500/j/Level-2-Memory-System/About-the-
L2-memory-system?lang=en
    [2] https://www.ti.com/lit/pdf/spruim2

    Fixes: 8abae9389bdb ("arm64: dts: ti: Add support for AM642 SoC")
    Reported-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Reviewed-by: Pratyush Yadav <p.yadav@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211113043635.4296-1-nm@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ada43142961e1bc67537d81067cb1571688de348
Author: Gaurav Jain <gaurav.jain@nxp.com>
Date:   Mon Nov 22 17:02:34 2021 +0530

    crypto: caam - save caam memory to support crypto engine retry mechanism.

    [ Upstream commit 087e1d715bccf25dc0e83294576e416b0386ba20 ]

    When caam queue is full (-ENOSPC), caam frees descriptor memory.
    crypto-engine checks if retry support is true and h/w queue
    is full(-ENOSPC), then requeue the crypto request.
    During processing the requested descriptor again, caam gives below error.
    (caam_jr 30902000.jr: 40000006: DECO: desc idx 0: Invalid KEY Command).

    This patch adds a check to return when caam input ring is full
    and retry support is true. so descriptor memory is not freed
    and requeued request can be processed again.

    Fixes: 2d653936eb2cf ("crypto: caam - enable crypto-engine retry mechanism")
    Signed-off-by: Gaurav Jain <gaurav.jain@nxp.com>
    Reviewed-by: Horia Geantă <horia.geanta@nxp.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ae8b70195ba637ac7a7c0ed31ced75e89da7c4b
Author: Ming Lei <ming.lei@redhat.com>
Date:   Fri Dec 3 10:39:35 2021 +0800

    null_blk: allow zero poll queues

    [ Upstream commit 2bfdbe8b7ebd17b5331071071a910fbabc64b436 ]

    There isn't any reason to not allow zero poll queues from user
    viewpoint.

    Also sometimes we need to compare io poll between poll mode and irq
    mode, so not allowing poll queues is bad.

    Fixes: 15dfc662ef31 ("null_blk: Fix handling of submit_queues and poll_queues
attributes")
    Cc: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
```

```
    Link: https://lore.kernel.org/r/20211203023935.3424042-1-ming.lei@redhat.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4e8f70e3e9cc9391f68e79be4477d3f9c89815f5
Author: Alexei Starovoitov <ast@kernel.org>
Date:   Wed Dec 1 10:10:34 2021 -0800

    libbpf: Clean gen_loader's attach kind.

    [ Upstream commit 19250f5fc0c283892a61f3abf9d65e6325f63897 ]

    The gen_loader has to clear attach_kind otherwise the programs
    without attach_btf_id will fail load if they follow programs
    with attach_btf_id.

    Fixes: 67234743736a ("libbpf: Generate loader program out of BPF ELF file.")
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20211201181040.23337-12-alexei.starovoitov@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b0c5be64016642a9365dcbbd5f64563afee6c6e2
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Wed Dec 1 23:13:10 2021 +0800

    drm/radeon/radeon_kms: Fix a NULL pointer dereference in radeon_driver_open_kms()

    [ Upstream commit ab50cb9df8896b39aae65c537a30de2c79c19735 ]

    In radeon_driver_open_kms(), radeon_vm_bo_add() is assigned to
    vm->ib_bo_va and passes and used in radeon_vm_bo_set_addr(). In
    radeon_vm_bo_set_addr(), there is a dereference of vm->ib_bo_va,
    which could lead to a NULL pointer dereference on failure of
    radeon_vm_bo_add().

    Fix this bug by adding a check of vm->ib_bo_va.

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_DRM_RADEON=m show no new warnings,
    and our static analyzer no longer warns about this code.

    Fixes: cc9e67e3d700 ("drm/radeon: fix VM IB handling")
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e9a93b7710650ec6be801b56073d87f50842d0cc
Author: Zhou Qingyang <zhou1615@umn.edu>
Date:   Fri Dec 3 00:17:36 2021 +0800

    drm/amdgpu: Fix a NULL pointer dereference in amdgpu_connector_lcd_native_mode()

    [ Upstream commit b220110e4cd442156f36e1d9b4914bb9e87b0d00 ]
```

```
    In amdgpu_connector_lcd_native_mode(), the return value of
    drm_mode_duplicate() is assigned to mode, and there is a dereference
    of it in amdgpu_connector_lcd_native_mode(), which will lead to a NULL
    pointer dereference on failure of drm_mode_duplicate().

    Fix this bug add a check of mode.

    This bug was found by a static analyzer. The analysis employs
    differential checking to identify inconsistent security operations
    (e.g., checks or kfrees) between two code paths and confirms that the
    inconsistent operations are not recovered in the current function or
    the callers, so they constitute bugs.

    Note that, as a bug found by static analysis, it can be a false
    positive or hard to trigger. Multiple researchers have cross-reviewed
    the bug.

    Builds with CONFIG_DRM_AMDGPU=m show no new warnings, and
    our static analyzer no longer warns about this code.

    Fixes: d38ceaf99ed0 ("drm/amdgpu: add core driver (v4)")
    Signed-off-by: Zhou Qingyang <zhou1615@umn.edu>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b0c6d22ccde9c5178e19eaf86e3bcba9345e4e7e
Author: Paul Gerber <Paul.Gerber@tq-group.com>
Date:   Mon Nov 22 12:42:25 2021 +0100

    thermal/drivers/imx8mm: Enable ADC when enabling monitor

    [ Upstream commit 3de89d8842a2b5d3dd22ebf97dd561ae0a330948 ]

    The i.MX 8MP has a ADC_PD bit in the TMU_TER register that controls the
    operating mode of the ADC:
    * 0 means normal operating mode
    * 1 means power down mode

    When enabling/disabling the TMU, the ADC operating mode must be set
    accordingly.

    i.MX 8M Mini & Nano are lacking this bit.

    Signed-off-by: Paul Gerber <Paul.Gerber@tq-group.com>
    Signed-off-by: Alexander Stein <alexander.stein@ew.tq-group.com>
    Fixes: 2b8f1f0337c5 ("thermal: imx8mm: Add i.MX8MP support")
    Link: https://lore.kernel.org/r/20211122114225.196280-1-alexander.stein@ew.tq-group.com
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 932f9026b9591cccfe2f9cb523351dc8c1d88eaf
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Sat Nov 27 17:08:36 2021 +0300

    pinctrl: mediatek: add a check for error in mtk_pinconf_bias_get_rsel()

    [ Upstream commit 9f9d17c228c89e38ed612500126daf626270be9a ]

    All the other mtk_hw_get_value() calls have a check for "if (err)" so
    we can add one here as well.  This silences a Smatch warning:

        drivers/pinctrl/mediatek/pinctrl-mtk-common-v2.c:819 mtk_pinconf_bias_get_rsel()
        error: uninitialized symbol 'pd'.
```

```
    Fixes: fb34a9ae383a ("pinctrl: mediatek: support rsel feature")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Link: https://lore.kernel.org/r/20211127140836.GB24002@kili
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eaa0728269f006ae111274f854233eb5bcf6e930
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Sat Nov 27 17:07:50 2021 +0300

    pinctrl: mediatek: uninitialized variable in mtk_pctrl_show_one_pin()

    [ Upstream commit 67bbbcb49b968a93251de7b23616d5aff5d3a726 ]

    The "try_all_type" variable is not set if (hw->soc->pull_type) is false
    leading to the following Smatch warning:

        drivers/pinctrl/mediatek/pinctrl-paris.c:599 mtk_pctrl_show_one_pin()
        error: uninitialized symbol 'try_all_type'.

    Fixes: fb34a9ae383a ("pinctrl: mediatek: support rsel feature")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Link: https://lore.kernel.org/r/20211127140750.GA24002@kili
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ad93d0ae1b9cf929cdd906d8512276cd0b06a25
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Tue Nov 23 19:36:51 2021 +0100

    ACPI: EC: Rework flushing of EC work while suspended to idle

    [ Upstream commit 4a9af6cac050dce2e895ec3205c4615383ad9112 ]

    The flushing of pending work in the EC driver uses drain_workqueue()
    to flush the event handling work that can requeue itself via
    advance_transaction(), but this is problematic, because that
    work may also be requeued from the query workqueue.

    Namely, if an EC transaction is carried out during the execution of
    a query handler, it involves calling advance_transaction() which
    may queue up the event handling work again.  This causes the kernel
    to complain about attempts to add a work item to the EC event
    workqueue while it is being drained and worst-case it may cause a
    valid event to be skipped.

    To avoid this problem, introduce two new counters, events_in_progress
    and queries_in_progress, incremented when a work item is queued on
    the event workqueue or the query workqueue, respectively, and
    decremented at the end of the corresponding work function, and make
    acpi_ec_dispatch_gpe() the workqueues in a loop until the both of
    these counters are zero (or system wakeup is pending) instead of
    calling acpi_ec_flush_work().

    At the same time, change __acpi_ec_flush_work() to call
    flush_workqueue() instead of drain_workqueue() to flush the event
    workqueue.

    While at it, use the observation that the work item queued in
    acpi_ec_query() cannot be pending at that time, because it is used
    only once, to simplify the code in there.

    Additionally, clean up a comment in acpi_ec_query() and adjust white
    space in acpi_ec_event_processor().
```

```
    Fixes: f0ac20c3f613 ("ACPI: EC: Fix flushing of pending work")
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 735f365005f32a032d24a0ed30b1085e6ed49ced
Author: William Kucharski <william.kucharski@oracle.com>
Date:   Wed Dec 1 09:56:58 2021 -0700

    cgroup: Trace event cgroup id fields should be u64

    [ Upstream commit e14da77113bb890d7bf9e5d17031bdd476a7ce5e ]

    Various trace event fields that store cgroup IDs were declared as
    ints, but cgroup_id(() returns a u64 and the structures and associated
    TP_printk() calls were not updated to reflect this.

    Fixes: 743210386c03 ("cgroup: use cgrp->kn->id as the cgroup ID")
    Signed-off-by: William Kucharski <william.kucharski@oracle.com>
    Reviewed-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3582ef39a834b696cf0a86d799c1265483c2a664
Author: Zack Rusin <zackr@vmware.com>
Date:   Fri Nov 5 15:38:44 2021 -0400

    drm/vmwgfx: Fail to initialize on broken configs

    [ Upstream commit c451af78f301ff5156998d571c37cab329c10051 ]

    Some of our hosts have a bug where rescaning a pci bus results in stale
    fifo memory being mapped on the host. This makes any fifo communication
    impossible resulting in various kernel crashes.

    Instead of unexpectedly crashing, predictably fail to load the driver
    which will preserve the system.

    Fixes: fb1d9738ca05 ("drm/vmwgfx: Add DRM driver for VMware Virtual GPU")
    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211105193845.258816-4-
zackr@vmware.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7aab7cc8d6e65505ba3e917610e874c048e28d2b
Author: Zack Rusin <zackr@vmware.com>
Date:   Fri Nov 5 15:38:42 2021 -0400

    drm/vmwgfx: Remove the deprecated lower mem limit

    [ Upstream commit 826c387d015247df396a91eadbaca94f0394853c ]

    TTM during the transition to the new page allocator lost the ability
    to constrain the allocations via the lower_mem_limit. The code has
    been unused since the change:
    256dd44bd897 ("drm/ttm: nuke old page allocator")
    and there's no reason to keep it.

    Fixes: 256dd44bd897 ("drm/ttm: nuke old page allocator")
    Signed-off-by: Zack Rusin <zackr@vmware.com>
    Reviewed-by: Martin Krastev <krastevm@vmware.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211105193845.258816-2-
zackr@vmware.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 5d0752563f07bc77369e1947f5a72c3372086059
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Dec 1 05:05:59 2021 +0300

    arm64: dts: qcom: msm8916: fix MMC controller aliases

    [ Upstream commit b0293c19d42f6d6951c2fab9a47fed50baf2c14d ]

    Change sdhcN aliases to mmcN to make them actually work. Currently the
    board uses non-standard aliases sdhcN, which do not work, resulting in
    mmc0 and mmc1 hosts randomly changing indices between boots.

    Fixes: c4da5a561627 ("arm64: dts: qcom: Add msm8916 sdhci configuration nodes")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20211201020559.1611890-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d960e4513aad39ade0872b0196bde267773a56b6
Author: Jun Miao <jun.miao@intel.com>
Date:   Tue Nov 16 07:23:02 2021 +0800

    rcu: Avoid alloc_pages() when recording stack

    [ Upstream commit 300c0c5e721834f484b03fa3062602dd8ff48413 ]

    The default kasan_record_aux_stack() calls stack_depot_save() with GFP_NOWAIT,
    which in turn can then call alloc_pages(GFP_NOWAIT, ...).  In general, however,
    it is not even possible to use either GFP_ATOMIC nor GFP_NOWAIT in certain
    non-preemptive contexts/RT kernel including raw_spin_locks (see gfp.h and
ab00db216c9c7).
    Fix it by instructing stackdepot to not expand stack storage via alloc_pages()
    in case it runs out by using kasan_record_aux_stack_noalloc().

    Jianwei Hu reported:
    BUG: sleeping function called from invalid context at kernel/locking/rtmutex.c:969
    in_atomic(): 0, irqs_disabled(): 1, non_block: 0, pid: 15319, name: python3
    INFO: lockdep is turned off.
    irq event stamp: 0
      hardirqs last  enabled at (0): [<0000000000000000>] 0x0
      hardirqs last disabled at (0): [<ffffffff856c8b13>] copy_process+0xaf3/0x2590
      softirqs last  enabled at (0): [<ffffffff856c8b13>] copy_process+0xaf3/0x2590
      softirqs last disabled at (0): [<0000000000000000>] 0x0
      CPU: 6 PID: 15319 Comm: python3 Tainted: G        W  O 5.15-rc7-preempt-rt #1
      Hardware name: Supermicro SYS-E300-9A-8C/A2SDi-8C-HLN4F, BIOS 1.1b 12/17/2018
      Call Trace:
        show_stack+0x52/0x58
        dump_stack+0xa1/0xd6
        ___might_sleep.cold+0x11c/0x12d
        rt_spin_lock+0x3f/0xc0
        rmqueue+0x100/0x1460
        rmqueue+0x100/0x1460
        mark_usage+0x1a0/0x1a0
        ftrace_graph_ret_addr+0x2a/0xb0
        rmqueue_pcplist.constprop.0+0x6a0/0x6a0
         __kasan_check_read+0x11/0x20
         __zone_watermark_ok+0x114/0x270
         get_page_from_freelist+0x148/0x630
         is_module_text_address+0x32/0xa0
         __alloc_pages_nodemask+0x2f6/0x790
         __alloc_pages_slowpath.constprop.0+0x12d0/0x12d0
         create_prof_cpu_mask+0x30/0x30
         alloc_pages_current+0xb1/0x150
         stack_depot_save+0x39f/0x490
         kasan_save_stack+0x42/0x50
```

```
          kasan_save_stack+0x23/0x50
          kasan_record_aux_stack+0xa9/0xc0
          __call_rcu+0xff/0x9c0
          call_rcu+0xe/0x10
          put_object+0x53/0x70
          __delete_object+0x7b/0x90
          kmemleak_free+0x46/0x70
          slab_free_freelist_hook+0xb4/0x160
          kfree+0xe5/0x420
          kfree_const+0x17/0x30
          kobject_cleanup+0xaa/0x230
          kobject_put+0x76/0x90
          netdev_queue_update_kobjects+0x17d/0x1f0
          ... ...
          ksys_write+0xd9/0x180
          __x64_sys_write+0x42/0x50
          do_syscall_64+0x38/0x50
          entry_SYSCALL_64_after_hwframe+0x44/0xa9

    Links:
https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/include/linux/kasa
n.h?id=7cb3007ce2da27ec02a1a3211941e7fe6875b642
    Fixes: 84109ab58590 ("rcu: Record kvfree_call_rcu() call stack for KASAN")
    Fixes: 26e760c9a7c8 ("rcu: kasan: record and print call_rcu() call stack")
    Reported-by: Jianwei Hu <jianwei.hu@windriver.com>
    Reviewed-by: Uladzislau Rezki (Sony) <urezki@gmail.com>
    Acked-by: Marco Elver <elver@google.com>
    Tested-by: Juri Lelli <juri.lelli@redhat.com>
    Signed-off-by: Jun Miao <jun.miao@intel.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed299bd66483280ad7ad9aedcd41d3fab8abf125
Author: Mark Rutland <mark.rutland@arm.com>
Date:   Mon Nov 29 13:06:51 2021 +0000

    powerpc: Avoid discarding flags in system_call_exception()

    [ Upstream commit 08b0af5b2affbe7419853e8dd1330e4b3fe27125 ]

    Some thread flags can be set remotely, and so even when IRQs are disabled,
    the flags can change under our feet. Thus, when setting flags we must use
    an atomic operation rather than a plain read-modify-write sequence, as a
    plain read-modify-write may discard flags which are concurrently set by a
    remote thread, e.g.

            // task A                       // task B
            tmp = A->thread_info.flags;
                                            set_tsk_thread_flag(A, NEWFLAG_B);
            tmp |= NEWFLAG_A;
            A->thread_info.flags = tmp;

    arch/powerpc/kernel/interrupt.c's system_call_exception() sets
    _TIF_RESTOREALL in the thread info flags with a read-modify-write, which
    may result in other flags being discarded.

    Elsewhere in the file it uses clear_bits() to atomically remove flag bits,
    so use set_bits() here for consistency with those.

    There may be reasons (e.g. instrumentation) that prevent the use of
    set_thread_flag() and clear_thread_flag() here, which would otherwise be
    preferable.

    Fixes: ae7aaecc3f2f78b7 ("powerpc/64s: system call rfscv workaround for TM bugs")
    Signed-off-by: Mark Rutland <mark.rutland@arm.com>
```

```
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Cc: Eirik Fuller <efuller@redhat.com>
    Cc: Michael Ellerman <mpe@ellerman.id.au>
    Cc: Nicholas Piggin <npiggin@gmail.com>
    Link: https://lore.kernel.org/r/20211129130653.2037928-10-mark.rutland@arm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec10895eccb748b44fbe6dda018aeffe1cd4f2ba
Author: Florian Westphal <fw@strlen.de>
Date:   Tue Nov 23 12:50:31 2021 +0100

    netfilter: bridge: add support for pppoe filtering

    [ Upstream commit 28b78ecffea8078d81466b2e01bb5a154509f1ba ]

    This makes 'bridge-nf-filter-pppoe-tagged' sysctl work for
    bridged traffic.

    Looking at the original commit it doesn't appear this ever worked:

     static unsigned int br_nf_post_routing(unsigned int hook, struct sk_buff **pskb,
    [..]
            if (skb->protocol == htons(ETH_P_8021Q)) {
                    skb_pull(skb, VLAN_HLEN);
                    skb->network_header += VLAN_HLEN;
    +       } else if (skb->protocol == htons(ETH_P_PPP_SES)) {
    +               skb_pull(skb, PPPOE_SES_HLEN);
    +               skb->network_header += PPPOE_SES_HLEN;
            }
     [..]
            NF_HOOK(... POST_ROUTING, ...)

    ... but the adjusted offsets are never restored.

    The alternative would be to rip this code out for good,
    but otoh we'd have to keep this anyway for the vlan handling
    (which works because vlan tag info is in the skb, not the packet
     payload).

    Reported-and-tested-by: Amish Chana <amish@3g.co.za>
    Fixes: 516299d2f5b6f97 ("[NETFILTER]: bridge-nf: filter bridged IPv4/IPv6 encapsulated
in pppoe traffic")
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35c1367b4749e8a4107d0ed2e4cbaed4ab34faba
Author: Jesper Dangaard Brouer <brouer@redhat.com>
Date:   Mon Nov 15 21:36:25 2021 +0100

    igc: AF_XDP zero-copy metadata adjust breaks SKBs on XDP_PASS

    [ Upstream commit 4fa8fcd3440101dbacf4fae91de69877ef751977 ]

    Driver already implicitly supports XDP metadata access in AF_XDP
    zero-copy mode, as xsk_buff_pool's xp_alloc() naturally set xdp_buff
    data_meta equal data.

    This works fine for XDP and AF_XDP, but if a BPF-prog adjust via
    bpf_xdp_adjust_meta() and choose to call XDP_PASS, then igc function
    igc_construct_skb_zc() will construct an invalid SKB packet. The
    function correctly include the xdp->data_meta area in the memcpy, but
    forgot to pull header to take metasize into account.

    Fixes: fc9df2a0b520 ("igc: Enable RX via AF_XDP zero-copy")
```

```
    Signed-off-by: Jesper Dangaard Brouer <brouer@redhat.com>
    Tested-by: Nechama Kraus <nechamax.kraus@linux.intel.com>
    Acked-by: Maciej Fijalkowski <maciej.fijalkowski@intel.com>
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1111c474adc451fcfa14edcda63eb84df2096e7
Author: Oleksij Rempel <linux@rempel-privat.de>
Date:   Wed Nov 17 11:34:26 2021 +0100

    thermal/drivers/imx: Implement runtime PM support

    [ Upstream commit 4cf2ddf16e175ee18c5c29865c32da7d6269cf44 ]

    Starting with commit d92ed2c9d3ff ("thermal: imx: Use driver's local
    data to decide whether to run a measurement") this driver stared using
    irq_enabled flag to make decision to power on/off the thermal
    core. This triggered a regression, where after reaching critical
    temperature, alarm IRQ handler set irq_enabled to false, disabled
    thermal core and was not able read temperature and disable cooling
    sequence.

    In case the cooling device is "CPU/GPU freq", the system will run with
    reduce performance until next reboot.

    To solve this issue, we need to move all parts implementing hand made
    runtime power management and let it handle actual runtime PM framework.

    Fixes: d92ed2c9d3ff ("thermal: imx: Use driver's local data to decide whether to run a
measurement")
    Signed-off-by: Oleksij Rempel <o.rempel@pengutronix.de>
    Tested-by: Petr Beneš <petr.benes@ysoft.com>
    Link: https://lore.kernel.org/r/20211117103426.81813-1-o.rempel@pengutronix.de
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f125c857e4f072a2c60c215d862a9c0a3eefb922
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Nov 30 11:38:37 2021 +0800

    net: lantiq: fix missing free_netdev() on error in ltq_etop_probe()

    [ Upstream commit 2680ce7fc9939221da16e86a2e73cc1df563c82c ]

    Add the missing free_netdev() before return from ltq_etop_probe()
    in the error handling case.

    Fixes: 14d4e308e0aa ("net: lantiq: configure the burst length in ethernet drivers")
    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 64825bc5e90cde865f9edd4fd59ceede7984a02e
Author: Bhupesh Sharma <bhupesh.sharma@linaro.org>
Date:   Mon Nov 29 01:28:54 2021 +0530

    net: stmmac: Add platform level debug register dump feature

    [ Upstream commit 4047b9db1aa7512a10ba3560a3f63821c8c40235 ]

    dwmac-qcom-ethqos currently exposes a mechanism to dump rgmii registers
    after the 'stmmac_dvr_probe()' returns. However with commit
    5ec55823438e ("net: stmmac: add clocks management for gmac driver"),
    we now let 'pm_runtime_put()' disable the clocks before returning from
    'stmmac_dvr_probe()'.
```

```
    This causes a crash when 'rgmii_dump()' register dumps are enabled,
    as the clocks are already off.

    Since other dwmac drivers (possible future users as well) might
    require a similar register dump feature, introduce a platform level
    callback to allow the same.

    This fixes the crash noticed while enabling rgmii_dump() dumps in
    dwmac-qcom-ethqos driver as well. It also allows future changes
    to keep a invoking the register dump callback from the correct
    place inside 'stmmac_dvr_probe()'.

    Fixes: 5ec55823438e ("net: stmmac: add clocks management for gmac driver")
    Cc: Joakim Zhang <qiangqing.zhang@nxp.com>
    Cc: David S. Miller <davem@davemloft.net>
    Signed-off-by: Bhupesh Sharma <bhupesh.sharma@linaro.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1968da1d4ad7e4ca5921b45a4144536828a47b6
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Thu Aug 19 22:05:28 2021 +0200

    media: venus: core: Fix a resource leak in the error handling path of 'venus_probe()'

    [ Upstream commit 8cc7a1b2aca067397a016cdb971a5e6ad9b640c7 ]

    A successful 'of_platform_populate()' call should be balanced by a
    corresponding 'of_platform_depopulate()' call in the error handling path
    of the probe, as already done in the remove function.

    A successful 'venus_firmware_init()' call should be balanced by a
    corresponding 'venus_firmware_deinit()' call in the error handling path
    of the probe, as already done in the remove function.

    Update the error handling path accordingly.

    Fixes: f9799fcce4bb ("media: venus: firmware: register separate platform_device for
firmware loader")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Stanimir Varbanov <stanimir.varbanov@linaro.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90f95923fd867b7e65eaad5ecab5119f489ea997
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Thu Aug 12 07:14:22 2021 +0200

    media: venus: core: Fix a potential NULL pointer dereference in an error handling path

    [ Upstream commit e4debea9be7d5db52bc6a565a4c02c3c6560d093 ]

    The normal path of the function makes the assumption that
    'pm_ops->core_power' may be NULL.
    We should make the same assumption in the error handling path or a NULL
    pointer dereference may occur.

    Add the missing test before calling 'pm_ops->core_power'

    Fixes: 9e8efdb57879 ("media: venus: core: vote for video-mem path")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Stanimir Varbanov <stanimir.varbanov@linaro.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit e5bb7f7c13b8a298f90e1b5f3a5e788db522c7fd
Author: Mansur Alisha Shaik <mansur@codeaurora.org>
Date:   Wed Oct 27 08:21:12 2021 +0200

    media: venus: correct low power frequency calculation for encoder

    [ Upstream commit b1f9bb8020783a48151e3a2864fbdc70548566dd ]

    In exististing implimentation, in min_loaded_core() for low_power
    vpp frequency value is considering as vpp_freq instead of low_power_freq.
    Fixed this by correcting vpp frequency calculation for encoder.

    Fixes: 3cfe5815ce0e (media: venus: Enable low power setting for encoder)
    Signed-off-by: Mansur Alisha Shaik <mansur@codeaurora.org>
    Signed-off-by: Stanimir Varbanov <stanimir.varbanov@linaro.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ba9e41ad20f6a99a43308c93c76250c20f2bc8f
Author: Philipp Zabel <p.zabel@pengutronix.de>
Date:   Fri Nov 19 11:41:20 2021 +0100

    media: coda: fix CODA960 JPEG encoder buffer overflow

    [ Upstream commit 1a59cd88f55068710f6549bee548846661673780 ]

    Stop the CODA960 JPEG encoder from overflowing capture buffers.
    The bitstream buffer overflow interrupt doesn't seem to be connected,
    so this has to be handled via timeout instead.

    Reported-by: Martin Weber <martin.weber@br-automation.com>
    Fixes: 96f6f62c4656 ("media: coda: jpeg: add CODA960 JPEG encoder support")
    Tested-by: Martin Weber <martin.weber@br-automation.com>
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c05a1394d4141e8590ffa371e43609a526bc1da
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Fri Nov 19 08:46:54 2021 +0100

    media: hantro: Hook up RK3399 JPEG encoder output

    [ Upstream commit 230d683ae04894933720425c8dead9508a09ebc3 ]

    The JPEG encoder found in the Hantro H1 encoder block only produces a
    raw entropy-encoded scan. The driver is responsible for building a JPEG
    compliant bitstream and placing the entropy-encoded scan in it. Right
    now the driver uses a bounce buffer for the hardware to output the raw
    scan to.

    In commit e765dba11ec2 ("hantro: Move hantro_enc_buf_finish to JPEG
    codec_ops.done"), the code that copies the raw scan from the bounce
    buffer to the capture buffer was moved, but was only hooked up for the
    Hantro H1 (then RK3288) variant. The RK3399 variant was broken,
    producing a JPEG bitstream without the scan, and the capture buffer's
    .bytesused field unset.

    Fix this by duplicating the code that is executed when the JPEG encoder
    finishes encoding a frame. As the encoded length is read back from
    hardware, and the variants having different register layouts, the
    code is duplicated rather than shared.
```

```
    Fixes: e765dba11ec2 ("hantro: Move hantro_enc_buf_finish to JPEG codec_ops.done")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Tested-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbcbd3aa68adecc536cc878c4a7216dce1d9ed4f
Author: Dafna Hirschfeld <dafna.hirschfeld@collabora.com>
Date:   Wed Nov 17 14:06:30 2021 +0100

    media: mtk-vcodec: call v4l2_m2m_ctx_release first when file is released

    [ Upstream commit 9f89c881bffbdffe4060ffaef3489a2830a6dd9c ]

    The func v4l2_m2m_ctx_release waits for currently running jobs
    to finish and then stop streaming both queues and frees the buffers.
    All this should be done before the call to mtk_vcodec_enc_release
    which frees the encoder handler. This fixes null-pointer dereference bug:

    [  638.028076] Mem abort info:
    [  638.030932]   ESR = 0x96000004
    [  638.033978]   EC = 0x25: DABT (current EL), IL = 32 bits
    [  638.039293]   SET = 0, FnV = 0
    [  638.042338]   EA = 0, S1PTW = 0
    [  638.045474]   FSC = 0x04: level 0 translation fault
    [  638.050349] Data abort info:
    [  638.053224]   ISV = 0, ISS = 0x00000004
    [  638.057055]   CM = 0, WnR = 0
    [  638.060018] user pgtable: 4k pages, 48-bit VAs, pgdp=000000012b6db000
    [  638.066485] [0000000000001a0] pgd=0000000000000000, p4d=0000000000000000
    [  638.073277] Internal error: Oops: 96000004 [#1] SMP
    [  638.078145] Modules linked in: rfkill mtk_vcodec_dec mtk_vcodec_enc uvcvideo mtk_mdp
mtk_vcodec_common videobuf2_dma_contig v4l2_h264 cdc_ether v4l2_mem2mem videobuf2_vmalloc
usbnet videobuf2_memops videobuf2_v4l2 r8152 videobuf2_common videodev cros_ec_sensors
cros_ec_sensors_core industrialio_triggered_buffer kfifo_buf elan_i2c elants_i2c sbs_battery
mc cros_usbpd_charger cros_ec_chardev cros_usbpd_logger crct10dif_ce mtk_vpu fuse ip_tables
x_tables ipv6
    [  638.118583] CPU: 0 PID: 212 Comm: kworker/u8:5 Not tainted 5.15.0-06427-
g58a1d4dcfc74-dirty #109
    [  638.127357] Hardware name: Google Elm (DT)
    [  638.131444] Workqueue: mtk-vcodec-enc mtk_venc_worker [mtk_vcodec_enc]
    [  638.137974] pstate: 60000005 (nZCv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--)
    [  638.144925] pc : vp8_enc_encode+0x34/0x2b0 [mtk_vcodec_enc]
    [  638.150493] lr : venc_if_encode+0xac/0x1b0 [mtk_vcodec_enc]
    [  638.156060] sp : ffff8000124d3c40
    [  638.159364] x29: ffff8000124d3c40 x28: 0000000000000000 x27: 0000000000000000
    [  638.166493] x26: 0000000000000000 x25: ffff0000e7f252d0 x24: ffff8000124d3d58
    [  638.173621] x23: ffff8000124d3d58 x22: ffff8000124d3d60 x21: 0000000000000001
    [  638.180750] x20: ffff80001137e000 x19: 0000000000000000 x18: 0000000000000001
    [  638.187878] x17: 000000040044ffff x16: 00400032b5503510 x15: 0000000000000000
    [  638.195006] x14: ffff8000118536c0 x13: ffff8000ee1da000 x12: 0000000030d4d91d
    [  638.202134] x11: 0000000000000000 x10: 0000000000000980 x9 : ffff8000124d3b20
    [  638.209262] x8 : ffff0000c18d4ea0 x7 : ffff0000c18d44c0 x6 : ffff0000c18d44c0
    [  638.216391] x5 : ffff80000904a3b0 x4 : ffff8000124d3d58 x3 : ffff8000124d3d60
    [  638.223519] x2 : ffff8000124d3d78 x1 : 0000000000000001 x0 : ffff80001137efb8
    [  638.230648] Call trace:
    [  638.233084]  vp8_enc_encode+0x34/0x2b0 [mtk_vcodec_enc]
    [  638.238304]  venc_if_encode+0xac/0x1b0 [mtk_vcodec_enc]
    [  638.243525]  mtk_venc_worker+0x110/0x250 [mtk_vcodec_enc]
    [  638.248918]  process_one_work+0x1f8/0x498
    [  638.252923]  worker_thread+0x140/0x538
    [  638.256664]  kthread+0x148/0x158
    [  638.259884]  ret_from_fork+0x10/0x20
    [  638.263455] Code: f90023f9 2a0103f5 aa0303f6 aa0403f8 (f940d277)
```

```
    [  638.269538] ---[ end trace e374fc10f8e181f5 ]---

    [gst-master] root@debian:~/gst-build# [  638.019193] Unable to handle kernel NULL
 pointer dereference at virtual address 00000000000001a0
    Fixes: 4e855a6efa547 ("[media] vcodec: mediatek: Add Mediatek V4L2 Video Encoder
 Driver")
    Signed-off-by: Dafna Hirschfeld <dafna.hirschfeld@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 9fc0817811b4e59369c0355438999b30f38457d3
 Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
 Date:   Thu Nov 11 14:17:51 2021 +0100

    media: mtk-vcodec: Fix an error handling path in 'mtk_vcodec_probe()'

    [ Upstream commit 615c6f28b9ad7efc9bfbef2cafc6a0c5bc0c21e0 ]

    In case of error the 'media_device_init()' call is not balanced by a
    corresponding 'media_device_cleanup()' call.

    Add it, when needed, as already done in the remove function.

    Fixes: 118add98f80e ("media: mtk-vcodec: vdec: add media device if using stateless api")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Acked-by: Tzung-Bi Shih <tzungbi@google.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 0bd798534cead04c1b84c6a10dc95e73b3becbd6
 Author: Yang Yingliang <yangyingliang@huawei.com>
 Date:   Fri Oct 15 11:58:55 2021 +0200

    media: si470x-i2c: fix possible memory leak in si470x_i2c_probe()

    [ Upstream commit ef054e345ed8c79ce1121a3599b5a2dfd78e57a0 ]

    n the 'radio->hdl.error' error handling, ctrl handler allocated by
    v4l2_ctrl_new_std() does not released, and caused memory leak as
    follows:

    unreferenced object 0xffff888033d54200 (size 256):
      comm "i2c-si470x-19", pid 909, jiffies 4294914203 (age 8.072s)
      hex dump (first 32 bytes):
        e8 69 11 03 80 88 ff ff 00 46 d5 33 80 88 ff ff  .i.......F.3....
        10 42 d5 33 80 88 ff ff 10 42 d5 33 80 88 ff ff  .B.3.....B.3....
      backtrace:
        [<00000000086bd4ed>] __kmalloc_node+0x1eb/0x360
        [<00000000bdb68871>] kvmalloc_node+0x66/0x120
        [<00000000fac74e4c>] v4l2_ctrl_new+0x7b9/0x1c60 [videodev]
        [<00000000693bf940>] v4l2_ctrl_new_std+0x19b/0x270 [videodev]
        [<00000000c0cb91bc>] si470x_i2c_probe+0x2d3/0x9a0 [radio_si470x_i2c]
        [<0000000056a6f01f>] i2c_device_probe+0x4d8/0xbe0

    Fix the error handling path to avoid memory leak.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Fixes: 8c081b6f9a9b ("media: radio: Critical v4l2 registration...")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 700189b49c16cb6b006c5fa8f41bf6f2fa6a5033
Author: Fabio Estevam <festevam@denx.de>
Date:   Fri Oct 8 15:10:14 2021 +0200

    media: imx-pxp: Initialize the spinlock prior to using it

    [ Upstream commit ed2f97ad4b21072f849cf4ae6645d1f2b1d3f550 ]

    After devm_request_threaded_irq() is called there is a chance that an
    interrupt may occur before the spinlock is initialized, which will trigger
    a kernel oops.

    To prevent that, move the initialization of the spinlock prior to
    requesting the interrupts.

    Fixes: 51abcf7fdb70 ("media: imx-pxp: add i.MX Pixel Pipeline driver")
    Signed-off-by: Fabio Estevam <festevam@denx.de>
    Reviewed-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 114ba26a2ed17c7c9451404d3a5dac1343a953bd
Author: Suresh Udipi <sudipi@jp.adit-jv.com>
Date:   Fri Aug 13 17:07:54 2021 +0200

    media: rcar-csi2: Correct the selection of hsfreqrange

    [ Upstream commit cee44d4fbacbbdfe62697ec94e76c6e4f726c5df ]

    hsfreqrange should be chosen based on the calculated mbps which
    is closer to the default bit rate  and within the range as per
    table[1]. But current calculation always selects first value which
    is greater than or equal to the calculated mbps which may lead
    to chosing a wrong range in some cases.

    For example for 360 mbps for H3/M3N
    Existing logic selects
    Calculated value 360Mbps : Default 400Mbps Range [368.125 -433.125 mbps]

    This hsfreqrange is out of range.

    The logic is changed to get the default value which is closest to the
    calculated value [1]

    Calculated value 360Mbps : Default 350Mbps  Range [320.625 -380.625 mpbs]

    [1] specs r19uh0105ej0200-r-car-3rd-generation.pdf [Table 25.9]

    Please note that According to Renesas in Table 25.9 the range for
    220 default value is corrected as below

     |Range (Mbps)     |  Default  Bit rate (Mbps) |
     ------------------------------------------------
     | 197.125-244.125 |     220                   |
     ------------------------------------------------

    Fixes: 769afd212b16 ("media: rcar-csi2: add Renesas R-Car MIPI CSI-2 receiver driver")
    Signed-off-by: Suresh Udipi <sudipi@jp.adit-jv.com>
    Signed-off-by: Kazuyoshi Akiyama <akiyama@nds-osk.co.jp>
    Signed-off-by: Michael Rodin <mrodin@de.adit-jv.com>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 4eeaf701cdba7c6db88ec3155f5ddcb992038670
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Tue Nov 23 01:00:02 2021 +0100

    media: i2c: ov8865: Fix lockdep error

    [ Upstream commit 6e1c9bc9ae96e57bcd8807174f2c0f44f9ef7938 ]

    ov8865_state_init() calls ov8865_state_mipi_configure() which uses
    __v4l2_ctrl_s_ctrl[_int64](). This means that sensor->mutex (which
    is also sensor->ctrls.handler.lock) must be locked before calling
    ov8865_state_init().

    Note ov8865_state_mipi_configure() is also used in other places where
    the lock is already held so it cannot be changed itself.

    This fixes the following lockdep kernel WARN:

    [   13.233421] WARNING: CPU: 0 PID: 8 at drivers/media/v4l2-core/v4l2-ctrls-api.c:833
__v4l2_ctrl_s_ctrl+0x4d/0x60 [videodev]
    ...
    [   13.234063] Call Trace:
    [   13.234074]  ov8865_state_configure+0x98b/0xc00 [ov8865]
    [   13.234095]  ov8865_probe+0x4b1/0x54c [ov8865]
    [   13.234117]  i2c_device_probe+0x13c/0x2d0

    Fixes: 11c0d8fdccc5 ("media: i2c: Add support for the OV8865 image sensor")
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33d9763e2cdb32aa06a856fd655dabebeb3cc824
Author: Daniel Scally <djrscally@gmail.com>
Date:   Tue Nov 23 01:00:01 2021 +0100

    media: i2c: Re-order runtime pm initialisation

    [ Upstream commit d2484fbf780762f6f9cc3abb7a07ee42dca2eaa3 ]

    The kerneldoc for pm_runtime_set_suspended() says:

    "It is not valid to call this function for devices with runtime PM
    enabled"

    To satisfy that requirement, re-order the calls so that
    pm_runtime_enable() is the last one.

    Fixes: 11c0d8fdccc5 ("media: i2c: Add support for the OV8865 image sensor")
    Signed-off-by: Daniel Scally <djrscally@gmail.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c247824072778c14804e3cb52c6163360af00a02
Author: Eugen Hristev <eugen.hristev@microchip.com>
Date:   Wed Nov 17 16:40:09 2021 +0100

    media: i2c: imx274: fix s_frame_interval runtime resume not requested

    [ Upstream commit da653498c20ba5b185214d8ae43b4e8e9594f520 ]

    pm_runtime_resume_and_get should be called when the s_frame_interval
    is called.
```

```
    The driver will try to access device registers to configure VMAX, coarse
    time and exposure.

    Currently if the runtime is not resumed, this fails:
     # media-ctl -d /dev/media0 --set-v4l2 '"IMX274 1-001a":0[fmt:SRGGB10_1X10/3840x2
    160@1/10]'

    IMX274 1-001a: imx274_binning_goodness: ask 3840x2160, size 3840x2160, goodness 0
    IMX274 1-001a: imx274_binning_goodness: ask 3840x2160, size 1920x1080, goodness -3000
    IMX274 1-001a: imx274_binning_goodness: ask 3840x2160, size 1280x720, goodness -4000
    IMX274 1-001a: imx274_binning_goodness: ask 3840x2160, size 1280x540, goodness -4180
    IMX274 1-001a: __imx274_change_compose: selected 1x1 binning
    IMX274 1-001a: imx274_set_frame_interval: input frame interval = 1 / 10
    IMX274 1-001a: imx274_read_mbreg : addr 0x300e, val=0x1 (2 bytes)
    IMX274 1-001a: imx274_set_frame_interval : register SVR = 1
    IMX274 1-001a: imx274_read_mbreg : addr 0x30f6, val=0x6a8 (2 bytes)
    IMX274 1-001a: imx274_set_frame_interval : register HMAX = 1704
    IMX274 1-001a: imx274_set_frame_length : input length = 2112
    IMX274 1-001a: imx274_write_mbreg : i2c bulk write failed, 30f8 = 884 (3 bytes)
    IMX274 1-001a: imx274_set_frame_length error = -121
    IMX274 1-001a: imx274_set_frame_interval error = -121
    Unable to setup formats: Remote I/O error (121)

    The device is not resumed thus the remote I/O error.

    Setting the frame interval works at streaming time, because
    pm_runtime_resume_and_get is called at s_stream time before sensor setup.
    The failure happens when only the s_frame_interval is called separately
    independently on streaming time.

    Fixes: ad97bc37426c ("media: i2c: imx274: Add IMX274 power on and off sequence")
    Signed-off-by: Eugen Hristev <eugen.hristev@microchip.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b96878b06c414047134afbda2e0508f67419c6c
Author: Alan Maguire <alan.maguire@oracle.com>
Date:   Mon Nov 29 10:00:40 2021 +0000

    libbpf: Silence uninitialized warning/error in btf_dump_dump_type_data

    [ Upstream commit 43174f0d4597325cb91f1f1f55263eb6e6101036 ]

    When compiling libbpf with gcc 4.8.5, we see:

      CC       staticobjs/btf_dump.o
    btf_dump.c: In function ‘btf_dump_dump_type_data.isra.24’:
    btf_dump.c:2296:5: error: ‘err’ may be used uninitialized in this function [-
Werror=maybe-uninitialized]
      if (err < 0)
         ^
    cc1: all warnings being treated as errors
    make: *** [staticobjs/btf_dump.o] Error 1

    While gcc 4.8.5 is too old to build the upstream kernel, it's possible it
    could be used to build standalone libbpf which suffers from the same problem.
    Silence the error by initializing 'err' to 0.  The warning/error seems to be
    a false positive since err is set early in the function.  Regardless we
    shouldn't prevent libbpf from building for this.

    Fixes: 920d16af9b42 ("libbpf: BTF dumper support for typed data")
    Signed-off-by: Alan Maguire <alan.maguire@oracle.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
```

```
    Link: https://lore.kernel.org/bpf/1638180040-8037-1-git-send-email-
alan.maguire@oracle.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1e3c38903ee5640158080fb8a35911697c3b3f38
Author: Jan Kara <jack@suse.cz>
Date:   Thu Nov 25 14:36:41 2021 +0100

    bfq: Do not let waker requests skip proper accounting

    [ Upstream commit c65e6fd460b4df796ecd6ea22e132076ed1f2820 ]

    Commit 7cc4ffc55564 ("block, bfq: put reqs of waker and woken in
    dispatch list") added a condition to bfq_insert_request() which added
    waker's requests directly to dispatch list. The rationale was that
    completing waker's IO is needed to get more IO for the current queue.
    Although this rationale is valid, there is a hole in it. The waker does
    not necessarily serve the IO only for the current queue and maybe it's
    current IO is not needed for current queue to make progress. Furthermore
    injecting IO like this completely bypasses any service accounting within
    bfq and thus we do not properly track how much service is waker's queue
    getting or that the waker is actually doing any IO. Depending on the
    conditions this can result in the waker getting too much or too few
    service.

    Consider for example the following job file:

    [global]
    directory=/mnt/repro/
    rw=write
    size=8g
    time_based
    runtime=30
    ramp_time=10
    blocksize=1m
    direct=0
    ioengine=sync

    [slowwriter]
    numjobs=1
    prioclass=2
    prio=7
    fsync=200

    [fastwriter]
    numjobs=1
    prioclass=2
    prio=0
    fsync=200

    Despite processes have very different IO priorities, they get the same
    about of service. The reason is that bfq identifies these processes as
    having waker-wakee relationship and once that happens, IO from
    fastwriter gets injected during slowwriter's time slice. As a result bfq
    is not aware that fastwriter has any IO to do and constantly schedules
    only slowwriter's queue. Thus fastwriter is forced to compete with
    slowwriter's IO all the time instead of getting its share of time based
    on IO priority.

    Drop the special injection condition from bfq_insert_request(). As a
    result, requests will be tracked and queued in a normal way and on next
    dispatch bfq_select_queue() can decide whether the waker's inserted
    requests should be injected during the current queue's timeslice or not.

    Fixes: 7cc4ffc55564 ("block, bfq: put reqs of waker and woken in dispatch list")
```

```
    Acked-by: Paolo Valente <paolo.valente@linaro.org>
    Signed-off-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20211125133645.27483-8-jack@suse.cz
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0907bda87fe72151520e748ce303582769ca943d
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Thu Oct 28 16:51:38 2021 +0300

    mfd: atmel-flexcom: Use .resume_noirq

    [ Upstream commit 5d051cf94fd5834a1513aa77e542c49fd973988a ]

    Flexcom IP embeds 3 other IPs: usart, i2c, spi and selects the operation
    mode (usart, i2c, spi) via mode register (FLEX_MR). On i2c bus there might
    be connected critical devices (like PMIC) which on suspend/resume should
    be suspended/resumed at the end/beginning. i2c uses
    .suspend_noirq/.resume_noirq for this kind of purposes. Align flexcom
    to use .resume_noirq as it should be resumed before the embedded IPs.
    Otherwise the embedded devices might behave badly.

    Fixes: 7fdec11015c3 ("atmel_flexcom: Support resuming after a chip reset")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Tested-by: Codrin Ciubotariu <codrin.ciubotariu@microchip.com>
    Acked-by: Nicolas Ferre <nicolas.ferre@microchip.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211028135138.3481166-3-claudiu.beznea@microchip.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df131ab1a4b647f29552b6639b0393227e2ecee8
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Thu Oct 28 16:51:37 2021 +0300

    mfd: atmel-flexcom: Remove #ifdef CONFIG_PM_SLEEP

    [ Upstream commit 8c0fad75dcaa650e3f3145a2c35847bc6a65cb7f ]

    Remove compilation flag and use __maybe_unused and pm_ptr instead.

    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Acked-by: Nicolas Ferre <nicolas.ferre@microchip.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Link: https://lore.kernel.org/r/20211028135138.3481166-2-claudiu.beznea@microchip.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dff83598dcdadfce94c97f09bb655bd5ca8ac361
Author: Alvin Šipraga <alsi@bang-olufsen.dk>
Date:   Mon Nov 29 11:30:19 2021 +0100

    net: dsa: rtl8365mb: set RGMII RX delay in steps of 0.3 ns

    [ Upstream commit ef136837aaf6f37f2dec78551671a6883f868d69 ]

    A contact at Realtek has clarified what exactly the units of RGMII RX
    delay are. The answer is that the unit of RX delay is "about 0.3 ns".
    Take this into account when parsing rx-internal-delay-ps by
    approximating the closest step value. Delays of more than 2.1 ns are
    rejected.

    This obviously contradicts the previous assumption in the driver that a
    step value of 4 was "about 2 ns", but Realtek also points out that it is
    easy to find more than one RX delay step value which makes RGMII work.

    Fixes: 4af2950c50c8 ("net: dsa: realtek-smi: add rtl8365mb subdriver for RTL8365MB-VC")
```

```
    Cc: Arınç ÜNAL <arinc.unal@arinc9.com>
    Signed-off-by: Alvin Šipraga <alsi@bang-olufsen.dk>
    Acked-by: Arınç ÜNAL <arinc.unal@arinc9.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5161d12c7344ce29c4f6036b3126216ed43593f
Author: Joey Gouly <joey.gouly@arm.com>
Date:   Sun Nov 21 16:56:42 2021 +0000

    pinctrl: apple: return an error if pinmux is missing in the DT

    [ Upstream commit 839930ca1bd0c79cdf370d11462ef4a81b664e44 ]

    If of_property_count_u32_elems returned 0, return -EINVAL to indicate
    a failure. Previously this would return 0.

    Fixes: a0f160ffcb83 ("pinctrl: add pinctrl/GPIO driver for Apple SoCs")
    Signed-off-by: Joey Gouly <joey.gouly@arm.com>
    Suggested-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Link: https://lore.kernel.org/r/20211121165642.27883-12-joey.gouly@arm.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a285fc361dcbcd0d95d606d05be90ae6177df68
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Thu Nov 25 11:00:18 2021 +0200

    tty: serial: atmel: Call dma_async_issue_pending()

    [ Upstream commit 4f4b9b5895614eb2e2b5f4cab7858f44bd113e1b ]

    The driver wrongly assummed that tx_submit() will start the transfer,
    which is not the case, now that the at_xdmac driver is fixed. tx_submit
    is supposed to push the current transaction descriptor to a pending queue,
    waiting for issue_pending to be called. issue_pending must start the
    transfer, not tx_submit.

    Fixes: 34df42f59a60 ("serial: at91: add rx dma support")
    Fixes: 08f738be88bb ("serial: at91: add tx dma support")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Link: https://lore.kernel.org/r/20211125090028.786832-4-tudor.ambarus@microchip.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1638bd5987e17c11ddb991c33c126a4524a38fb9
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Thu Nov 25 11:00:17 2021 +0200

    tty: serial: atmel: Check return code of dmaengine_submit()

    [ Upstream commit 1e67bd2b8cb90b66e89562598e9c2046246832d3 ]

    The tx_submit() method of struct dma_async_tx_descriptor is entitled
    to do sanity checks and return errors if encountered. It's not the
    case for the DMA controller drivers that this client is using
    (at_h/xdmac), because they currently don't do sanity checks and always
    return a positive cookie at tx_submit() method. In case the controller
    drivers will implement sanity checks and return errors, print a message
    so that the client will be informed that something went wrong at
    tx_submit() level.

    Fixes: 08f738be88bb ("serial: at91: add tx dma support")
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Acked-by: Richard Genoud <richard.genoud@gmail.com>
```

```
    Link: https://lore.kernel.org/r/20211125090028.786832-3-tudor.ambarus@microchip.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71499ad0a25435f16438fb923cf386877b965c6e
Author: Peng Fan <peng.fan@nxp.com>
Date:   Fri Nov 12 14:31:55 2021 +0800

    arm64: dts: ti: k3-j721e: correct cache-sets info

    [ Upstream commit 7a0df1f969c14939f60a7f9a6af72adcc314675f ]

    A72 Cluster has 48KB Icache, 32KB Dcache and 1MB L2 Cache
     - ICache is 3-way set-associative
     - Dcache is 2-way set-associative
     - Line size are 64bytes

    So correct the cache-sets info.

    Fixes: 2d87061e70dea ("arm64: dts: ti: Add Support for J721E SoC")
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Link: https://lore.kernel.org/r/20211112063155.3485777-1-peng.fan@oss.nxp.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2457654a3e7b29522b8660b9d01d9fa3d87ea6cc
Author: Anilkumar Kolli <akolli@codeaurora.org>
Date:   Wed Nov 24 19:11:31 2021 +0200

    ath11k: Use host CE parameters for CE interrupts configuration

    [ Upstream commit b689f091aafd1a874b2f88137934276ab0fca480 ]

    CE interrupt configuration uses host ce parameters to assign/free
    interrupts. Use host ce parameters to enable/disable interrupts.
    This patch fixes below BUG,

    BUG: KASAN: global-out-of-bounds in 0xffffffbffdfb035c at addr
    ffffffbffde6eeac
     Read of size 4 by task kworker/u8:2/132
     Address belongs to variable ath11k_core_qmi_firmware_ready+0x1b0/0x5bc [ath11k]

    OOB is due to ath11k_ahb_ce_irqs_enable() iterates ce_count(which is 12)
    times and accessing 12th element in target_ce_config
    (which has only 11 elements) from ath11k_ahb_ce_irq_enable().

    With this change host ce configs are used to enable/disable interrupts.

    Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-00471-QCAHKSWPL_SILICONZ-1

    Fixes: 967c1d1131fa ("ath11k: move target ce configs to hw_params")
    Signed-off-by: Anilkumar Kolli <akolli@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/1637249558-12793-1-git-send-email-akolli@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a25c6d65cbd0e48ee4c424c5f0e50edb07fc7c0f
Author: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
Date:   Wed Nov 17 14:30:35 2021 +0000

    crypto: qat - fix undetected PFVF timeout in ACK loop

    [ Upstream commit 5002200b4fedd7e90e4fbc2e5c42a4b3351df814 ]
```

```
    If the remote function did not ACK the reception of a message, the
    function __adf_iov_putmsg() could detect it as a collision.

    This was due to the fact that the collision and the timeout checks after
    the ACK loop were in the wrong order. The timeout must be checked at the
    end of the loop, so fix by swapping the order of the two checks.

    Fixes: 9b768e8a3909 ("crypto: qat - detect PFVF collision after ACK")
    Signed-off-by: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
    Co-developed-by: Marco Chiappero <marco.chiappero@intel.com>
    Signed-off-by: Marco Chiappero <marco.chiappero@intel.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 52e948654337ff32db7713f5069d597f3312d998
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Tue Nov 23 16:23:18 2021 -0800

    libbpf: Fix using invalidated memory in bpf_linker

    [ Upstream commit 593835377f24ca1bb98008ec1dc3baefe491ad6e ]

    add_dst_sec() can invalidate bpf_linker's section index making
    dst_symtab pointer pointing into unallocated memory. Reinitialize
    dst_symtab pointer on each iteration to make sure it's always valid.

    Fixes: faf6ed321cf6 ("libbpf: Add BPF static linker APIs")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211124002325.1737739-7-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4485f09a46dab6b0482ca3dd4b4fc1fe0906cf0c
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Tue Nov 23 16:23:17 2021 -0800

    libbpf: Fix glob_syms memory leak in bpf_linker

    [ Upstream commit 8cb125566c40b7141d8842c534f0ea5820ee3d5c ]

    glob_syms array wasn't freed on bpf_link__free(). Fix that.

    Fixes: a46349227cd8 ("libbpf: Add linker extern resolution support for functions and
global variables")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211124002325.1737739-6-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05e7801383c3b8b252b59a20e88618c5b420ef2b
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Tue Nov 23 16:23:14 2021 -0800

    libbpf: Fix potential misaligned memory access in btf_ext__new()

    [ Upstream commit 401891a9debaf0a684502f2aaecf53448cee9414 ]

    Perform a memory copy before we do the sanity checks of btf_ext_hdr.
    This prevents misaligned memory access if raw btf_ext data is not 4-byte
    aligned ([0]).

    While at it, also add missing const qualifier.

      [0] Closes: https://github.com/libbpf/libbpf/issues/391
```

```
    Fixes: 2993e0515bb4 ("tools/bpf: add support to read .BTF.ext sections")
    Reported-by: Evgeny Vereshchagin <evvers@ya.ru>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211124002325.1737739-3-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 72c6bd95b5bbb63f5663429a8355427221cca0e4
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Tue Nov 23 16:23:13 2021 -0800

    tools/resolve_btf_ids: Close ELF file on error

    [ Upstream commit 1144ab9bdf3430e1b5b3f22741e5283841951add ]

    Fix one case where we don't do explicit clean up.

    Fixes: fbbb68de80a4 ("bpf: Add resolve_btfids tool to resolve BTF IDs in ELF object")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211124002325.1737739-2-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 395d283799d6ca75498b11953c1c94f659ebfb2f
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Tue Nov 23 12:01:04 2021 -0800

    libbpf: Load global data maps lazily on legacy kernels

    [ Upstream commit 16e0c35c6f7a2e90d52f3035ecf942af21417b7b ]

    Load global data maps lazily, if kernel is too old to support global
    data. Make sure that programs are still correct by detecting if any of
    the to-be-loaded programs have relocation against any of such maps.

    This allows to solve the issue ([0]) with bpf_printk() and Clang
    generating unnecessary and unreferenced .rodata.strX.Y sections, but it
    also goes further along the CO-RE lines, allowing to have a BPF object
    in which some code can work on very old kernels and relies only on BPF
    maps explicitly, while other BPF programs might enjoy global variable
    support. If such programs are correctly set to not load at runtime on
    old kernels, bpf_object will load and function correctly now.

      [0] https://lore.kernel.org/bpf/CAK-
59YFPU3qO+_pXWOH+c1LSA=8WA1yabJZfREjOEXNHAqgXNg@mail.gmail.com/

    Fixes: aed659170a31 ("libbpf: Support multiple .rodata.* and .data.* BPF maps")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Song Liu <songliubraving@fb.com>
    Link: https://lore.kernel.org/bpf/20211123200105.387855-1-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 039803713975df27811755eb946813eabc3667e9
Author: Dillon Min <dillon.minfei@gmail.com>
Date:   Sat Jul 24 11:44:02 2021 +0800

    ARM: dts: stm32: fix dtbs_check warning on ili9341 dts binding on stm32f429 disco

    [ Upstream commit b046049e59dca5e5830dc75ed16acf7657a95161 ]

    Since the compatible string defined from ilitek,ili9341.yaml is
    "st,sf-tc240t-9370-t", "ilitek,ili9341"

    so, append "ilitek,ili9341" to avoid the below dtbs_check warning.
```

```
    arch/arm/boot/dts/stm32f429-disco.dt.yaml: display@1: compatible:
    ['st,sf-tc240t-9370-t'] is too short

    Fixes: a726e2f000ec ("ARM: dts: stm32: enable ltdc binding with ili9341, gyro l3gd20 on
stm32429-disco board")
    Signed-off-by: Dillon Min <dillon.minfei@gmail.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9911ac27e96facd8dd2c8842e29aadc9859c80fe
Author: Vladimir Zapolskiy <vladimir.zapolskiy@linaro.org>
Date:   Thu Nov 11 17:48:07 2021 +0200

    cpufreq: qcom-hw: Fix probable nested interrupt handling

    [ Upstream commit e0e27c3d4e20dab861566f1c348ae44e4b498630 ]

    Re-enabling an interrupt from its own interrupt handler may cause
    an interrupt storm, if there is a pending interrupt and because its
    handling is disabled due to already done entrance into the handler
    above in the stack.

    Also, apparently it is improper to lock a mutex in an interrupt contex.

    Fixes: 275157b367f4 ("cpufreq: qcom-cpufreq-hw: Add dcvs interrupt support")
    Signed-off-by: Vladimir Zapolskiy <vladimir.zapolskiy@linaro.org>
    Reviewed-by: Matthias Kaehlcke <mka@chromium.org>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a938ea3a0da03eab3dfef2c11eee55beda99f1c1
Author: Adam Ford <aford173@gmail.com>
Date:   Sat Nov 20 13:39:16 2021 -0600

    soc: imx: gpcv2: keep i.MX8MM VPU-H1 bus clock active

    [ Upstream commit 8361b8b29f9389084b679db854cf733375c64763 ]

    Enable the vpu-h1 clock when the domain is active because reading
    or writing to the VPU-H1 IP block cause the system to hang.

    Fixes: 656ade7aa42a ("soc: imx: gpcv2: keep i.MX8M* bus clocks enabled")
    Signed-off-by: Adam Ford <aford173@gmail.com>
    Reviewed-by: Fabio Estevam <festevam@gmail.com>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 966a596b5ea24614d29a7a853bab9d7261671c8e
Author: Lukasz Luba <lukasz.luba@arm.com>
Date:   Tue Nov 9 19:57:12 2021 +0000

    cpufreq: qcom-cpufreq-hw: Update offline CPUs per-cpu thermal pressure

    [ Upstream commit 93d9e6f93e1586fcc97498c764be2e8c8401f4bd ]

    The thermal pressure signal gives information to the scheduler about
    reduced CPU capacity due to thermal. It is based on a value stored in
    a per-cpu 'thermal_pressure' variable. The online CPUs will get the
    new value there, while the offline won't. Unfortunately, when the CPU
    is back online, the value read from per-cpu variable might be wrong
    (stale data).  This might affect the scheduler decisions, since it
```

```
    sees the CPU capacity differently than what is actually available.

    Fix it by making sure that all online+offline CPUs would get the
    proper value in their per-cpu variable when there is throttling
    or throttling is removed.

    Fixes: 275157b367f479 ("cpufreq: qcom-cpufreq-hw: Add dcvs interrupt support")
    Reviewed-by: Thara Gopinath <thara.gopinath@linaro.org>
    Signed-off-by: Lukasz Luba <lukasz.luba@arm.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit abb7dee4674e853af2fa67b8732d015223cfec47
Author: George G. Davis <davis.george@siemens.com>
Date:   Fri Jul 16 16:49:35 2021 -0400

    mtd: hyperbus: rpc-if: fix bug in rpcif_hb_remove

    [ Upstream commit baaf965f94308301d2dc554d72a87d7432cd5ce6 ]

    The following KASAN BUG is observed when testing the rpc-if driver on
    rcar-gen3:

    root@rcar-gen3:~# modprobe -r rpc-if
    [  101.930146] ==================================================================
    [  101.937408] BUG: KASAN: slab-out-of-bounds in __lock_acquire+0x518/0x25d0
    [  101.944240] Read of size 8 at addr ffff0004c5be2750 by task modprobe/664
    [  101.950959]
    [  101.952466] CPU: 2 PID: 664 Comm: modprobe Not tainted 5.14.0-rc1-00342-g1a1464d7aa31
#1
    [  101.960578] Hardware name: Renesas H3ULCB board based on r8a77951 (DT)
    [  101.967120] Call trace:
    [  101.969580]  dump_backtrace+0x0/0x2c0
    [  101.973275]  show_stack+0x1c/0x30
    [  101.976616]  dump_stack_lvl+0x9c/0xd8
    [  101.980301]  print_address_description.constprop.0+0x74/0x2b8
    [  101.986071]  kasan_report+0x1f4/0x26c
    [  101.989757]  __asan_load8+0x98/0xd4
    [  101.993266]  __lock_acquire+0x518/0x25d0
    [  101.997215]  lock_acquire.part.0+0x18c/0x360
    [  102.001506]  lock_acquire+0x74/0x90
    [  102.005013]  _raw_spin_lock_irq+0x98/0x130
    [  102.009131]  __pm_runtime_disable+0x30/0x210
    [  102.013427]  rpcif_hb_remove+0x5c/0x70 [rpc_if]
    [  102.018001]  platform_remove+0x40/0x80
    [  102.021771]  __device_release_driver+0x234/0x350
    [  102.026412]  driver_detach+0x158/0x20c
    [  102.030179]  bus_remove_driver+0xa0/0x140
    [  102.034212]  driver_unregister+0x48/0x80
    [  102.038153]  platform_driver_unregister+0x18/0x24
    [  102.042879]  rpcif_platform_driver_exit+0x1c/0x34 [rpc_if]
    [  102.048400]  __arm64_sys_delete_module+0x210/0x310
    [  102.053212]  invoke_syscall+0x60/0x190
    [  102.056986]  el0_svc_common+0x12c/0x144
    [  102.060844]  do_el0_svc+0x88/0xac
    [  102.064181]  el0_svc+0x24/0x3c
    [  102.067257]  el0t_64_sync_handler+0x1a8/0x1b0
    [  102.071634]  el0t_64_sync+0x198/0x19c
    [  102.075315]
    [  102.076815] Allocated by task 628:
    [  102.080781]
    [  102.082280] Last potentially related work creation:
    [  102.087524]
    [  102.089022] The buggy address belongs to the object at ffff0004c5be2000
    [  102.089022]  which belongs to the cache kmalloc-2k of size 2048
```

```
     [  102.101555] The buggy address is located 1872 bytes inside of
     [  102.101555]  2048-byte region [ffff0004c5be2000, ffff0004c5be2800)
     [  102.113486] The buggy address belongs to the page:
     [  102.118409]
     [  102.119908] Memory state around the buggy address:
     [  102.124711]  ffff0004c5be2600: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     [  102.131947]  ffff0004c5be2680: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     [  102.139181] >ffff0004c5be2700: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     [  102.146412]                                                 ^
     [  102.152257]  ffff0004c5be2780: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     [  102.159491]  ffff0004c5be2800: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc
     [  102.166723] ==================================================================

     The above bug is caused by use of the wrong pointer in the
     rpcif_disable_rpm() call. Fix the bug by using the correct pointer.

     Fixes: 5de15b610f78 ("mtd: hyperbus: add Renesas RPC-IF driver")
     Signed-off-by: George G. Davis <davis.george@siemens.com>
     Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
     Link: https://lore.kernel.org/r/20210716204935.25859-1-george_davis@mentor.com
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit d886e0a5156d7b34ce209b65b46f8c7d10201165
  Author: Prasad Malisetty <pmaliset@codeaurora.org>
  Date:   Tue Nov 16 16:31:48 2021 +0530

     arm64: dts: qcom: sc7280: Fix 'interrupt-map' parent address cells

     [ Upstream commit 66b788133030f0c69a0ecc7f72f7939b119c9a69 ]

     Update interrupt-map parent address cells for sc7280
     Similar to existing Qcom SoCs.

     Fixes: 92e0ee9f8 ("arm64: dts: qcom: sc7280: Add PCIe and PHY related nodes")
     Signed-off-by: Prasad Malisetty <pmaliset@codeaurora.org>
     Reviewed-by: Stephen Boyd <swboyd@chromium.org>
     Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
     Link: https://lore.kernel.org/r/1637060508-30375-4-git-send-email-
  pmaliset@codeaurora.org
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 2010518b9179051e85234c79e10d553af8fc88c6
  Author: Prasad Malisetty <pmaliset@codeaurora.org>
  Date:   Tue Nov 16 16:31:46 2021 +0530

     arm64: dts: qcom: sc7280: Fix incorrect clock name

     [ Upstream commit fa09b2248714c64644576d8064e9bd292a504a0e ]

     Replace pcie_1_pipe-clk clock name with pcie_1_pipe_clk
     To match with dt binding.

     Fixes: ab7772de8612 ("arm64: dts: qcom: SC7280: Add rpmhcc clock controller node")
     Signed-off-by: Prasad Malisetty <pmaliset@codeaurora.org>
     Reviewed-by: Stephen Boyd <swboyd@chromium.org>
     Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
     Link: https://lore.kernel.org/r/1637060508-30375-2-git-send-email-
  pmaliset@codeaurora.org
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit d4af631b955dc7976b4a8cb445cdf0e9ba0b794c
  Author: Chengfeng Ye <cyeaa@connect.ust.hk>
  Date:   Thu Nov 4 06:46:42 2021 -0700

     crypto: qce - fix uaf on qce_skcipher_register_one
```

```
    [ Upstream commit e9c195aaeed1b45c9012adbe29dedb6031e85aa8 ]

    Pointer alg points to sub field of tmpl, it
    is dereferenced after tmpl is freed. Fix
    this by accessing alg before free tmpl.

    Fixes: ec8f5d8f ("crypto: qce - Qualcomm crypto engine driver")
    Signed-off-by: Chengfeng Ye <cyeaa@connect.ust.hk>
    Acked-by: Thara Gopinath <thara.gopinath@linaro.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aca83fc60399babe243a93fdf7708fd3bf74ded4
Author: Chengfeng Ye <cyeaa@connect.ust.hk>
Date:   Thu Nov 4 06:38:31 2021 -0700

    crypto: qce - fix uaf on qce_ahash_register_one

    [ Upstream commit b4cb4d31631912842eb7dce02b4350cbb7562d5e ]

    Pointer base points to sub field of tmpl, it
    is dereferenced after tmpl is freed. Fix
    this by accessing base before free tmpl.

    Fixes: ec8f5d8f ("crypto: qce - Qualcomm crypto engine driver")
    Signed-off-by: Chengfeng Ye <cyeaa@connect.ust.hk>
    Acked-by: Thara Gopinath <thara.gopinath@linaro.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b77b9e7f32cd6561907f987ab45a720379c3da91
Author: Chengfeng Ye <cyeaa@connect.ust.hk>
Date:   Thu Nov 4 06:28:07 2021 -0700

    crypto: qce - fix uaf on qce_aead_register_one

    [ Upstream commit 4a9dbd021970ffe1b92521328377b699acba7c52 ]

    Pointer alg points to sub field of tmpl, it
    is dereferenced after tmpl is freed. Fix
    this by accessing alg before free tmpl.

    Fixes: 9363efb4 ("crypto: qce - Add support for AEAD algorithms")
    Signed-off-by: Chengfeng Ye <cyeaa@connect.ust.hk>
    Acked-by: Thara Gopinath <thara.gopinath@linaro.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef890bb693e22e865b41bfd0c8a4242f56da2009
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Mon Nov 1 14:02:33 2021 +0000

    crypto: keembay-ocs-ecc - Fix error return code in kmb_ocs_ecc_probe()

    [ Upstream commit 94ad2d19a97efdb603a21fcad0625f466f1cdd0f ]

    Fix to return negative error code -ENOMEM from the error handling
    case instead of 0, as done elsewhere in this function.

    Fixes: c9f608c38009 ("crypto: keembay-ocs-ecc - Add Keem Bay OCS ECC Driver")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Reviewed-by: Daniele Alessandrelli <daniele.alessandrelli@intel.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
```

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

commit ef414d72b8a5a724588e252665e16f1d68e172c0
Author: Tudor Ambarus <tudor.ambarus@microchip.com>
Date:   Wed Oct 27 16:30:01 2021 +0300

    crypto: atmel-aes - Reestablish the correct tfm context at dequeue

    [ Upstream commit 6d48de655917a9d782953eba65de4e3db593ddf0 ]

    In case there were more requests from different tfms in the crypto
    queue, only the context of the last initialized tfm was considered.

    Fixes: ec2088b66f7a ("crypto: atmel-aes - Allocate aes dev at tfm init time")
    Reported-by: Wolfgang Ocker <weo@reccoware.de>
    Signed-off-by: Tudor Ambarus <tudor.ambarus@microchip.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3d1a76c12ff83a324cb1448da30517f0020085e8
Author: Wang Hai <wanghai38@huawei.com>
Date:   Fri Oct 15 16:57:41 2021 +0800

    media: dmxdev: fix UAF when dvb_register_device() fails

    [ Upstream commit ab599eb11882f834951c436cc080c3455ba32b9b ]

    I got a use-after-free report:

    dvbdev: dvb_register_device: failed to create device dvb1.dvr0 (-12)
    ...
    ==================================================================
    BUG: KASAN: use-after-free in dvb_dmxdev_release+0xce/0x2f0
    ...
    Call Trace:
     dump_stack_lvl+0x6c/0x8b
     print_address_description.constprop.0+0x48/0x70
     kasan_report.cold+0x82/0xdb
     __asan_load4+0x6b/0x90
     dvb_dmxdev_release+0xce/0x2f0
    ...
    Allocated by task 7666:
     kasan_save_stack+0x23/0x50
     __kasan_kmalloc+0x83/0xa0
     kmem_cache_alloc_trace+0x22e/0x470
     dvb_register_device+0x12f/0x980
     dvb_dmxdev_init+0x1f3/0x230
    ...
    Freed by task 7666:
     kasan_save_stack+0x23/0x50
     kasan_set_track+0x20/0x30
     kasan_set_free_info+0x24/0x40
     __kasan_slab_free+0xf2/0x130
     kfree+0xd1/0x5c0
     dvb_register_device.cold+0x1ac/0x1fa
     dvb_dmxdev_init+0x1f3/0x230
    ...

    When dvb_register_device() in dvb_dmxdev_init() fails, dvb_dmxdev_init()
    does not return a failure, and the memory pointed to by dvbdev or
    dvr_dvbdev is invalid at this point. If they are used subsequently, it
    will result in UFA or null-ptr-deref.

    If dvb_register_device() in dvb_dmxdev_init() fails, fix the bug by making
    dvb_dmxdev_init() return an error as well.
```

```
    Link: https://lore.kernel.org/linux-media/20211015085741.1203283-1-wanghai38@huawei.com

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Hai <wanghai38@huawei.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 68ec6c28c635a59d2ff2e3198275b73a6c872cc3
Author: Biju Das <biju.das.jz@bp.renesas.com>
Date:   Mon Nov 15 14:28:30 2021 +0000

    arm64: dts: renesas: cat875: Add rx/tx delays

    [ Upstream commit e1a9faddffe7e555304dc2e3284c84fbee0679ee ]

    The CAT875 sub board from Silicon Linux uses a Realtek PHY.

    The phy driver commit bbc4d71d63549bcd003 ("net: phy: realtek: fix
    rtl8211e rx/tx delay config") introduced NFS mount failures.  Now it
    needs both rx/tx delays for the NFS mount to work.

    This patch fixes the NFS mount failure issue by adding "rgmii-id" mode
    to the avb device node.

    Signed-off-by: Biju Das <biju.das.jz@bp.renesas.com>
    Fixes: bbc4d71d63549bcd ("net: phy: realtek: fix rtl8211e rx/tx delay config")
    Link: https://lore.kernel.org/r/20211115142830.12651-1-biju.das.jz@bp.renesas.com
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aaac97c8d49530322bb5aec01328f5db990defb7
Author: Seevalamuthu Mariappan <quic_seevalam@quicinc.com>
Date:   Wed Nov 17 09:39:41 2021 +0200

    ath11k: add hw_param for wakeup_mhi

    [ Upstream commit 081e2d6476e30399433b509684d5da4d1844e430 ]

    Wakeup mhi is needed before pci_read/write only for QCA6390 and WCN6855. Since
    wakeup & release mhi is enabled for all hardwares, below mhi assert is seen in
    QCN9074 when doing 'rmmod ath11k_pci':

            Kernel panic - not syncing: dev_wake != 0
            CPU: 2 PID: 13535 Comm: procd Not tainted 4.4.60 #1
            Hardware name: Generic DT based system
            [<80316dac>] (unwind_backtrace) from [<80313700>] (show_stack+0x10/0x14)
            [<80313700>] (show_stack) from [<805135dc>] (dump_stack+0x7c/0x9c)
            [<805135dc>] (dump_stack) from [<8032136c>] (panic+0x84/0x1f8)
            [<8032136c>] (panic) from [<80549b24>] (mhi_pm_disable_transition+0x3b8/0x5b8)
            [<80549b24>] (mhi_pm_disable_transition) from [<80549ddc>]
(mhi_power_down+0xb8/0x100)
            [<80549ddc>] (mhi_power_down) from [<7f5242b0>]
(ath11k_mhi_op_status_cb+0x284/0x3ac [ath11k_pci])
            [E][__mhi_device_get_sync] Did not enter M0 state, cur_state:RESET
pm_state:SHUTDOWN Process
            [E][__mhi_device_get_sync] Did not enter M0 state, cur_state:RESET
pm_state:SHUTDOWN Process
            [E][__mhi_device_get_sync] Did not enter M0 state, cur_state:RESET
pm_state:SHUTDOWN Process
            [<7f5242b0>] (ath11k_mhi_op_status_cb [ath11k_pci]) from [<7f524878>]
(ath11k_mhi_stop+0x10/0x20 [ath11k_pci])
            [<7f524878>] (ath11k_mhi_stop [ath11k_pci]) from [<7f525b94>]
(ath11k_pci_power_down+0x54/0x90 [ath11k_pci])
```

```
             [<7f525b94>] (ath11k_pci_power_down [ath11k_pci]) from [<8056b2a8>]
  (pci_device_shutdown+0x30/0x44)
             [<8056b2a8>] (pci_device_shutdown) from [<805cfa0c>]
  (device_shutdown+0x124/0x174)
             [<805cfa0c>] (device_shutdown) from [<8033aaa4>] (kernel_restart+0xc/0x50)
             [<8033aaa4>] (kernel_restart) from [<8033ada8>] (SyS_reboot+0x178/0x1ec)
             [<8033ada8>] (SyS_reboot) from [<80301b80>] (ret_fast_syscall+0x0/0x34)

    Hence, disable wakeup/release mhi using hw_param for other hardwares.

    Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.5.0.1-01060-QCAHKSWPL_SILICONZ-1

    Fixes: a05bd8513335 ("ath11k: read and write registers below unwindowed address")
    Signed-off-by: Seevalamuthu Mariappan <quic_seevalam@quicinc.com>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/1636702019-26142-1-git-send-email-
  quic_seevalam@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit cb15c1b2b17e0331b2f8e0b4c292da11cabb30be
  Author: P Praneesh <ppranees@codeaurora.org>
  Date:   Fri Nov 12 11:01:26 2021 +0200

    ath11k: allocate dst ring descriptors from cacheable memory

    [ Upstream commit 6452f0a3d5651bb7edfd9c709e78973aaa4d3bfc ]

    tcl_data and reo_dst rings are currently being allocated using
    dma_allocate_coherent() which is non cacheable.

    Allocating ring memory from cacheable memory area allows cached descriptor
    access and prefetch next descriptors to optimize CPU usage during
    descriptor processing on NAPI. Based on the hardware param we can enable
    or disable this feature for the corresponding platform.

    Tested-on: QCN9074 hw1.0 PCI WLAN.HK.2.4.0.1.r2-00012-QCAHKSWPL_SILICONZ-1
    Tested-on: IPQ8074 hw2.0 AHB WLAN.HK.2.4.0.1-01695-QCAHKSWPL_SILICONZ-1

    Co-developed-by: Pradeep Kumar Chitrapu <pradeepc@codeaurora.org>
    Signed-off-by: Pradeep Kumar Chitrapu <pradeepc@codeaurora.org>
    Co-developed-by: Sriram R <srirrama@codeaurora.org>
    Signed-off-by: Sriram R <srirrama@codeaurora.org>
    Signed-off-by: Jouni Malinen <jouni@codeaurora.org>
    Signed-off-by: P Praneesh <ppranees@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/1630560820-21905-3-git-send-email-
  ppranees@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 30326fca2588775bf5225b92392cb92fe9981b0b
  Author: Wen Gong <wgong@codeaurora.org>
  Date:   Thu Oct 28 10:46:28 2021 +0300

    ath11k: set correct NL80211_FEATURE_DYNAMIC_SMPS for WCN6855

    [ Upstream commit 82c434c103408842a87404e873992b7698b6df2b ]

    Commit 6f4d70308e5e ("ath11k: support SMPS configuration for 6 GHz") changed
    "if (ht_cap & WMI_HT_CAP_DYNAMIC_SMPS)" to "if (ht_cap &
    WMI_HT_CAP_DYNAMIC_SMPS || ar->supports_6ghz)" which means
    NL80211_FEATURE_DYNAMIC_SMPS is enabled for all chips which support 6 GHz.
    However, WCN6855 supports 6 GHz but it does not support feature
    NL80211_FEATURE_DYNAMIC_SMPS, and this can lead to MU-MIMO test failures for
    WCN6855.
```

```
    Disable NL80211_FEATURE_DYNAMIC_SMPS for WCN6855 since its ht_cap does not
    support WMI_HT_CAP_DYNAMIC_SMPS. Enable the feature only on QCN9074 as that's
    the only other device supporting 6 GHz band.

    Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1

    Signed-off-by: Wen Gong <wgong@codeaurora.org>
    Signed-off-by: Jouni Malinen <jouni@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210914163726.38604-3-jouni@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc8d6f7e71bca1664b36d2d7e989fe3e0ca53fc3
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Thu Nov 18 14:12:33 2021 +0300

    drm/vboxvideo: fix a NULL vs IS_ERR() check

    [ Upstream commit cebbb5c46d0cb0615fd0c62dea9b44273d0a9780 ]

    The devm_gen_pool_create() function never returns NULL, it returns
    error pointers.

    Fixes: 4cc9b565454b ("drm/vboxvideo: Use devm_gen_pool_create")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211118111233.GA1147@kili
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de886282a2f4d7648d2878f1ab552cb4adb6aa77
Author: Jeremy Kerr <jk@codeconstruct.com.au>
Date:   Thu Nov 18 14:57:23 2021 +0800

    mctp/test: Update refcount checking in route fragment tests

    [ Upstream commit f6ef47e5bdc6f652176e433b02317fc83049f8d7 ]

    In 99ce45d5e, we moved a route refcount decrement from
    mctp_do_fragment_route into the caller. This invalidates the assumption
    that the route test makes about refcount behaviour, so the route tests
    fail.

    This change fixes the test case to suit the new refcount behaviour.

    Fixes: 99ce45d5e7db ("mctp: Implement extended addressing")
    Signed-off-by: Jeremy Kerr <jk@codeconstruct.com.au>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16fd41dc5ef944dabb350ccce30283e601066518
Author: Tirthendu Sarkar <tirthendu.sarkar@intel.com>
Date:   Wed Nov 17 18:06:13 2021 +0530

    selftests/bpf: Fix xdpxceiver failures for no hugepages

    [ Upstream commit dd7f091fd22b1dce6c20e8f7769aa068ed88ac6d ]

    xsk_configure_umem() needs hugepages to work in unaligned mode. So when
    hugepages are not configured, 'unaligned' tests should be skipped which
    is determined by the helper function hugepages_present(). This function
    erroneously returns true with MAP_NORESERVE flag even when no hugepages
    are configured. The removal of this flag fixes the issue.

    The test TEST_TYPE_UNALIGNED_INV_DESC also needs to be skipped when
```

```
    there are no hugepages. However, this was not skipped as there was no
    check for presence of hugepages and hence was failing. The check to skip
    the test has now been added.

    Fixes: a4ba98dd0c69 (selftests: xsk: Add test for unaligned mode)
    Signed-off-by: Tirthendu Sarkar <tirthendu.sarkar@intel.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20211117123613.22288-1-tirthendu.sarkar@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cae61f6e4059336460f4c2c83d54fe8492267536
Author: Lyude Paul <lyude@redhat.com>
Date:   Fri Nov 5 14:33:40 2021 -0400

    drm/dp: Don't read back backlight mode in drm_edp_backlight_enable()

    [ Upstream commit 646596485e1ed2182adf293dfd5aec4a96c46330 ]

    As it turns out, apparently some machines will actually leave additional
    backlight functionality like dynamic backlight control on before the OS
    loads. Currently we don't take care to disable unsupported features when
    writing back the backlight mode, which can lead to some rather strange
    looking behavior when adjusting the backlight.

    So, let's fix this by just not reading back the current backlight mode on
    initial enable. I don't think there should really be any downsides to this,
    and this will ensure we don't leave any unsupported functionality enabled.

    This should fix at least one (but not all) of the issues seen with DPCD
    backlight support on fi-bdw-samus

    v5:
    * Just avoid reading back DPCD register - Doug Anderson

    Signed-off-by: Lyude Paul <lyude@redhat.com>
    Fixes: 867cf9cd73c3 ("drm/dp: Extract i915's eDP backlight code into DRM helpers")
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211105183342.130810-4-
lyude@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a845ca01e51f1d9076686e7daf5c1a167b16b7f5
Author: Alexander Aring <aahringo@redhat.com>
Date:   Wed Nov 17 09:20:43 2021 -0500

    fs: dlm: fix build with CONFIG_IPV6 disabled

    [ Upstream commit 1b9beda83e27a0c2cd75d1cb743c297c7b36c844 ]

    This patch will surround the AF_INET6 case in sk_error_report() of dlm
    with a #if IS_ENABLED(CONFIG_IPV6). The field sk->sk_v6_daddr is not
    defined when CONFIG_IPV6 is disabled. If CONFIG_IPV6 is disabled, the
    socket creation with AF_INET6 should already fail because a runtime
    check if AF_INET6 is registered. However if there is the possibility
    that AF_INET6 is set as sk_family the sk_error_report() callback will
    print then an invalid family type error.

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: 4c3d90570bcc ("fs: dlm: don't call kernel_getpeername() in error_report()")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5b88bc7e35bda1c64b045822e01121c96c9b4a0
Author: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
```

```
Date:   Thu Nov 4 13:27:06 2021 +0100

    kernel/locking: Use a pointer in ww_mutex_trylock().

    [ Upstream commit 2202e15b2b1a946ce760d96748cd7477589701ab ]

    mutex_acquire_nest() expects a pointer, pass the pointer.

    Fixes: 12235da8c80a1 ("kernel/locking: Add context to ww_mutex_trylock()")
    Signed-off-by: Sebastian Andrzej Siewior <bigeasy@linutronix.de>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lkml.kernel.org/r/20211104122706.frk52zxbjorso2kv@linutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 034e93a7e9b0bf3838352ca02c0b9c6ea11ea295
Author: Jens Wiklander <jens.wiklander@linaro.org>
Date:   Tue Jun 15 22:23:50 2021 +0200

    tee: fix put order in teedev_close_context()

    [ Upstream commit f18397ab3ae23e8e43bba9986e66af6d4497f2ad ]

    Prior to this patch was teedev_close_context() calling tee_device_put()
    before teedev_ctx_put() leading to teedev_ctx_release() accessing
    ctx->teedev just after the reference counter was decreased on the
    teedev. Fix this by calling teedev_ctx_put() before tee_device_put().

    Fixes: 217e0250cccb ("tee: use reference counting for tee_context")
    Reviewed-by: Sumit Garg <sumit.garg@linaro.org>
    Signed-off-by: Jens Wiklander <jens.wiklander@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 704569e394a521e21424667f02c7900faa05d6a7
Author: oujiefeng <oujiefeng@huawei.com>
Date:   Wed Nov 17 09:21:19 2021 +0800

    spi: hisi-kunpeng: Fix the debugfs directory name incorrect

    [ Upstream commit 40fafc8eca3f0d41b9dade5c10afb2dad723aad7 ]

    Change the debugfs directory name from hisi_spi65535 to hisi_spi0.

    Fixes: 2b2142f247eb ("spi: hisi-kunpeng: Add debugfs support")
    Signed-off-by: oujiefeng <oujiefeng@huawei.com>
    Signed-off-by: Jay Fang <f.fangjian@huawei.com>
    Link: https://lore.kernel.org/r/20211117012119.55558-1-f.fangjian@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac8cd936e9b90a644bea6d5fb4d9b000538152d2
Author: Karthikeyan Kathirvel <kathirve@codeaurora.org>
Date:   Mon Nov 15 11:04:41 2021 +0100

    ath11k: reset RSN/WPA present state for open BSS

    [ Upstream commit 64bc3aa02ae78b1fcb1b850e0eb1f0622002bfaa ]

    The ath11k driver is caching the information about RSN/WPA IE in the
    configured beacon template. The cached information is used during
    associations to figure out whether 4-way PKT/2-way GTK peer flags need to
    be set or not.

    But the code never cleared the state when no such IE was found. This can
    for example happen when moving from an WPA/RSN to an open setup. The
    (seemingly connected) peer was then not able to communicate over the
```

```
    link because the firmware assumed a different (encryption enabled) state
    for the peer.

    Tested-on: IPQ6018 hw1.0 AHB WLAN.HK.2.5.0.1-01100-QCAHKSWPL_SILICONZ-1

    Fixes: 01e34233c645 ("ath11k: fix wmi peer flags in peer assoc command")
    Cc: Venkateswara Naralasetty <vnaralas@codeaurora.org>
    Reported-by: Sven Eckelmann <sven@narfation.org>
    Signed-off-by: Karthikeyan Kathirvel <kathirve@codeaurora.org>
    [sven@narfation.org: split into separate patches, clean up commit message]
    Signed-off-by: Sven Eckelmann <sven@narfation.org>

    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211115100441.33771-2-sven@narfation.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e21fbc2f230035e5dc02451a4306e85cb3fe60e
Author: Karthikeyan Kathirvel <kathirve@codeaurora.org>
Date:   Mon Nov 15 11:04:40 2021 +0100

    ath11k: clear the keys properly via DISABLE_KEY

    [ Upstream commit 436a4e88659842a7cf634d7cc088c8f2cc94ebf5 ]

    DISABLE_KEY sets the key_len to 0, firmware will not delete the keys if
    key_len is 0. Changing from security mode to open mode will cause mcast
    to be still encrypted without vdev restart.

    Set the proper key_len for DISABLE_KEY cmd to clear the keys in
    firmware.

    Tested-on: IPQ6018 hw1.0 AHB WLAN.HK.2.5.0.1-01100-QCAHKSWPL_SILICONZ-1

    Fixes: d5c65159f289 ("ath11k: driver for Qualcomm IEEE 802.11ax devices")
    Reported-by: Sven Eckelmann <sven@narfation.org>
    Signed-off-by: Karthikeyan Kathirvel <kathirve@codeaurora.org>
    [sven@narfation.org: split into separate patches, clean up commit message]
    Signed-off-by: Sven Eckelmann <sven@narfation.org>

    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211115100441.33771-1-sven@narfation.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33abd70ac278e19a41a0cfe0e3636d434674ea43
Author: Sven Eckelmann <sven@narfation.org>
Date:   Mon Nov 15 11:29:55 2021 +0200

    ath11k: Fix ETSI regd with weather radar overlap

    [ Upstream commit 086c921a354089f209318501038d43c98d3f409f ]

    Some ETSI countries have a small overlap in the wireless-regdb with an ETSI
    channel (5590-5650). A good example is Australia:

      country AU: DFS-ETSI
            (2400 - 2483.5 @ 40), (36)
            (5150 - 5250 @ 80), (23), NO-OUTDOOR, AUTO-BW
            (5250 - 5350 @ 80), (20), NO-OUTDOOR, AUTO-BW, DFS
            (5470 - 5600 @ 80), (27), DFS
            (5650 - 5730 @ 80), (27), DFS
            (5730 - 5850 @ 80), (36)
            (57000 - 66000 @ 2160), (43), NO-OUTDOOR

    If the firmware (or the BDF) is shipped with these rules then there is only
    a 10 MHz overlap with the weather radar:
```

```
    * below: 5470 - 5590
    * weather radar: 5590 - 5600
    * above: (none for the rule "5470 - 5600 @ 80")

    There are several wrong assumption in the ath11k code:

    * there is always a valid range below the weather radar
      (actually: there could be no range below the weather radar range OR range
       could be smaller than 20 MHz)
    * intersected range in the weather radar range is valid
      (actually: the range could be smaller than 20 MHz)
    * range above weather radar is either empty or valid
      (actually: the range could be smaller than 20 MHz)

    These wrong assumption will lead in this example to a rule

      (5590 - 5600 @ 20), (N/A, 27), (600000 ms), DFS, AUTO-BW

    which is invalid according to is_valid_reg_rule() because the freq_diff is
    only 10 MHz but the max_bandwidth is set to 20 MHz. Which results in a
    rejection like:

      WARNING: at backports-20210222_001-4.4.60-b157d2276/net/wireless/reg.c:3984
      [...]
      Call trace:
      [<ffffffbffc3d2e50>] reg_get_max_bandwidth+0x300/0x3a8 [cfg80211]
      [<ffffffbffc3d3d0c>] regulatory_set_wiphy_regd_sync+0x3c/0x98 [cfg80211]
      [<ffffffbffc651598>] ath11k_regd_update+0x1a8/0x210 [ath11k]
      [<ffffffbffc652108>] ath11k_regd_update_work+0x18/0x20 [ath11k]
      [<ffffffc0000a93e0>] process_one_work+0x1f8/0x340
      [<ffffffc0000a9784>] worker_thread+0x25c/0x448
      [<ffffffc0000aedc8>] kthread+0xd0/0xd8
      [<ffffffc000085550>] ret_from_fork+0x10/0x40
      ath11k c000000.wifi: failed to perform regd update : -22
      Invalid regulatory domain detected

    To avoid this, the algorithm has to be changed slightly. Instead of
    splitting a rule which overlaps with the weather radar range into 3 pieces
    and accepting the first two parts blindly, it must actually be checked for
    each piece whether it is a valid range. And only if it is valid, add it to
    the output array.

    When these checks are in place, the processed rules for AU would end up as

      country AU: DFS-ETSI
              (2400 - 2483 @ 40), (N/A, 36), (N/A)
              (5150 - 5250 @ 80), (6, 23), (N/A), NO-OUTDOOR, AUTO-BW
              (5250 - 5350 @ 80), (6, 20), (0 ms), NO-OUTDOOR, DFS, AUTO-BW
              (5470 - 5590 @ 80), (6, 27), (0 ms), DFS, AUTO-BW
              (5650 - 5730 @ 80), (6, 27), (0 ms), DFS, AUTO-BW
              (5730 - 5850 @ 80), (6, 36), (N/A), AUTO-BW

    and will be accepted by the wireless regulatory code.

    Fixes: d5c65159f289 ("ath11k: driver for Qualcomm IEEE 802.11ax devices")
    Signed-off-by: Sven Eckelmann <sven@narfation.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211112153116.1214421-1-sven@narfation.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1da1c9d3f198d1c7e37b382e005f65d06809e9e5
Author: Jackie Liu <liuyun01@kylinos.cn>
Date:   Tue Nov 16 09:17:17 2021 +0800
```

```
    Bluetooth: fix uninitialized variables notify_evt

    [ Upstream commit a27c519a816437ec92f0ffa3adbc168c2c08725b ]

    Coverity Scan report:

    [...]
    *** CID 1493985:  Uninitialized variables  (UNINIT)
    /net/bluetooth/hci_event.c: 4535 in hci_sync_conn_complete_evt()
    4529
    4530            /* Notify only in case of SCO over HCI transport data path which
    4531             * is zero and non-zero value shall be non-HCI transport data path
    4532             */
    4533            if (conn->codec.data_path == 0) {
    4534                    if (hdev->notify)
    >>>     CID 1493985:  Uninitialized variables  (UNINIT)
    >>>     Using uninitialized value "notify_evt" when calling "*hdev->notify".
    4535                            hdev->notify(hdev, notify_evt);
    4536            }
    4537
    4538            hci_connect_cfm(conn, ev->status);
    4539            if (ev->status)
    4540                    hci_conn_del(conn);
    [...]

    Although only btusb uses air_mode, and he only handles HCI_NOTIFY_ENABLE_SCO_CVSD
    and HCI_NOTIFY_ENABLE_SCO_TRANSP, there is still a very small chance that
    ev->air_mode is not equal to 0x2 and 0x3, but notify_evt is initialized to
    HCI_NOTIFY_ENABLE_SCO_CVSD or HCI_NOTIFY_ENABLE_SCO_TRANSP. the context is
    maybe not correct.

    Let us directly use the required function instead of re-initializing it,
    so as to restore the original logic and make the code more correct.

    Addresses-Coverity: ("Uninitialized variables")
    Fixes: f4f9fa0c07bb ("Bluetooth: Allow usb to auto-suspend when SCO use non-HCI
transport")
    Suggested-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Jackie Liu <liuyun01@kylinos.cn>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 640a476e38fce6c70c738ac09e5167e34f570303
Author: Pavel Skripkin <paskripkin@gmail.com>
Date:   Mon Nov 1 10:12:12 2021 +0300

    Bluetooth: stop proccessing malicious adv data

    [ Upstream commit 3a56ef719f0b9682afb8a86d64b2399e36faa4e6 ]

    Syzbot reported slab-out-of-bounds read in hci_le_adv_report_evt(). The
    problem was in missing validaion check.

    We should check if data is not malicious and we can read next data block.
    If we won't check ptr validness, code can read a way beyond skb->end and
    it can cause problems, of course.

    Fixes: e95beb414168 ("Bluetooth: hci_le_adv_report_evt code refactoring")
    Reported-and-tested-by: syzbot+e3fcb9c4f3c2a931dc40@syzkaller.appspotmail.com
    Signed-off-by: Pavel Skripkin <paskripkin@gmail.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 27f401fd974475205ca15f5e653a91f748730c26
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
```

```
Date:   Mon Oct 25 21:56:29 2021 +0100

    memory: renesas-rpc-if: Return error in case devm_ioremap_resource() fails

    [ Upstream commit 818fdfa89baac77a8df5a2c30f4fb798cc937aa0 ]

    Make sure we return error in case devm_ioremap_resource() fails for dirmap
    resource.

    Fixes: ca7d8b980b67 ("memory: add Renesas RPC-IF driver")
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Reviewed-by: Biju Das <biju.das.jz@bp.renesas.com>
    Reviewed-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20211025205631.21151-6-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b28bf312a273eb57272f1b79d4f5380a3a34d327
Author: Alexander Aring <aahringo@redhat.com>
Date:   Mon Nov 15 08:57:05 2021 -0500

    fs: dlm: don't call kernel_getpeername() in error_report()

    [ Upstream commit 4c3d90570bcc2b338f70f61f01110268e281ca3c ]

    In some cases kernel_getpeername() will held the socket lock which is
    already held when the socket layer calls error_report() callback. Since
    commit 9dfc685e0262 ("inet: remove races in inet{6}_getname()") this
    problem becomes more likely because the socket lock will be held always.
    You will see something like:

    bob9-u5 login: [  562.316860] BUG: spinlock recursion on CPU#7, swapper/7/0
    [  562.318562]  lock: 0xffff8f2284720088, .magic: dead4ead, .owner: swapper/7/0,
.owner_cpu: 7
    [  562.319522] CPU: 7 PID: 0 Comm: swapper/7 Not tainted 5.15.0+ #135
    [  562.320346] Hardware name: Red Hat KVM/RHEL-AV, BIOS 1.13.0-
2.module+el8.3.0+7353+9de0a3cc 04/01/2014
    [  562.321277] Call Trace:
    [  562.321529]  <IRQ>
    [  562.321734]  dump_stack_lvl+0x33/0x42
    [  562.322282]  do_raw_spin_lock+0x8b/0xc0
    [  562.322674]  lock_sock_nested+0x1e/0x50
    [  562.323057]  inet_getname+0x39/0x110
    [  562.323425]  ? sock_def_readable+0x80/0x80
    [  562.323838]  lowcomms_error_report+0x63/0x260 [dlm]
    [  562.324338]  ? wait_for_completion_interruptible_timeout+0xd2/0x120
    [  562.324949]  ? lock_timer_base+0x67/0x80
    [  562.325330]  ? do_raw_spin_unlock+0x49/0xc0
    [  562.325735]  ? _raw_spin_unlock_irqrestore+0x1e/0x40
    [  562.326218]  ? del_timer+0x54/0x80
    [  562.326549]  sk_error_report+0x12/0x70
    [  562.326919]  tcp_validate_incoming+0x3c8/0x530
    [  562.327347]  ? kvm_clock_read+0x14/0x30
    [  562.327718]  ? ktime_get+0x3b/0xa0
    [  562.328055]  tcp_rcv_established+0x121/0x660
    [  562.328466]  tcp_v4_do_rcv+0x132/0x260
    [  562.328835]  tcp_v4_rcv+0xcea/0xe20
    [  562.329173]  ip_protocol_deliver_rcu+0x35/0x1f0
    [  562.329615]  ip_local_deliver_finish+0x54/0x60
    [  562.330050]  ip_local_deliver+0xf7/0x110
    [  562.330431]  ? inet_rtm_getroute+0x211/0x840
    [  562.330848]  ? ip_protocol_deliver_rcu+0x1f0/0x1f0
    [  562.331310]  ip_rcv+0xe1/0xf0
```

```
[  562.331603]  ? ip_local_deliver+0x110/0x110
[  562.332011]  __netif_receive_skb_core+0x46a/0x1040
[  562.332476]  ? inet_gro_receive+0x263/0x2e0
[  562.332885]  __netif_receive_skb_list_core+0x13b/0x2c0
[  562.333383]  netif_receive_skb_list_internal+0x1c8/0x2f0
[  562.333896]  ? update_load_avg+0x7e/0x5e0
[  562.334285]  gro_normal_list.part.149+0x19/0x40
[  562.334722]  napi_complete_done+0x67/0x160
[  562.335134]  virtnet_poll+0x2ad/0x408 [virtio_net]
[  562.335644]  __napi_poll+0x28/0x140
[  562.336012]  net_rx_action+0x23d/0x300
[  562.336414]  __do_softirq+0xf2/0x2ea
[  562.336803]  irq_exit_rcu+0xc1/0xf0
[  562.337173]  common_interrupt+0xb9/0xd0

    It is and was always forbidden to call kernel_getpeername() in context
    of error_report(). To get rid of the problem we access the destination
    address for the peer over the socket structure. While on it we fix to
    print out the destination port of the inet socket.

    Fixes: 1a31833d085a ("DLM: Replace nodeid_to_addr with kernel_getpeername")
    Reported-by: Bob Peterson <rpeterso@redhat.com>
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4bd22ec63810c1e330f02d35e343e0dfb06ee63a
Author: Christian Hewitt <christianshewitt@gmail.com>
Date:   Tue Oct 12 05:25:21 2021 +0000

    arm64: dts: meson-gxbb-wetek: fix missing GPIO binding

    [ Upstream commit c019abb2feba3cbbd7cf7178f8e6499c4fa6fced ]

    The absence of this binding appears to be harmless in Linux but it breaks
    Ethernet support in mainline u-boot. So add the binding (which is present
    in all other u-boot supported GXBB device-trees).

    Fixes: fb72c03e0e32 ("ARM64: dts: meson-gxbb-wetek: add a wetek specific dtsi to cleanup
hub and play2")

    Signed-off-by: Christian Hewitt <christianshewitt@gmail.com>
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://lore.kernel.org/r/20211012052522.30873-3-christianshewitt@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f055b3d9656c9830b802e2bbf28314b572f9469
Author: Christian Hewitt <christianshewitt@gmail.com>
Date:   Tue Oct 12 05:25:20 2021 +0000

    arm64: dts: meson-gxbb-wetek: fix HDMI in early boot

    [ Upstream commit 8182a35868db5f053111d5d9d4da8fcb3f99259d ]

    Mark the VDDIO_AO18 regulator always-on and set hdmi-supply for the hdmi_tx
    node to ensure HDMI is powered in the early stages of boot.

    Fixes: fb72c03e0e32 ("ARM64: dts: meson-gxbb-wetek: add a wetek specific dtsi to cleanup
hub and play2")

    Signed-off-by: Christian Hewitt <christianshewitt@gmail.com>
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://lore.kernel.org/r/20211012052522.30873-2-christianshewitt@gmail.com
```

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

commit 3638abb5836ec8590c94a9718cb24e33b9bb4cee
Author: Alexander Stein <alexander.stein@mailbox.org>
Date:   Tue Oct 26 20:28:13 2021 +0200

    arm64: dts: amlogic: Fix SPI NOR flash node name for ODROID N2/N2+

    [ Upstream commit 95d35256b564aca33fb661eac77dc94bfcffc8df ]

    Fix the schema warning: "spi-flash@0: $nodename:0: 'spi-flash@0' does
     not match '^flash(@.*)?$'" from jedec,spi-nor.yaml

    Fixes: a084eaf3096c ("arm64: dts: meson-g12b-odroid-n2: add SPIFC controller node")
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Alexander Stein <alexander.stein@mailbox.org>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://lore.kernel.org/r/20211026182813.900775-3-alexander.stein@mailbox.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8de2d5c92c7c03447cbf76da24992be8d1c5d9a7
Author: Alexander Stein <alexander.stein@mailbox.org>
Date:   Tue Oct 26 20:28:12 2021 +0200

    arm64: dts: amlogic: meson-g12: Fix GPU operating point table node name

    [ Upstream commit bb98a6fd0b0e227cefb2ba91cea2b55455f203b7 ]

    Starting with commit 94274f20f6bf ("dt-bindings: opp: Convert to DT
    schema") the opp node name has a mandatory pattern. This change
    fixes the dtbs_check warning:
    gpu-opp-table: $nodename:0: 'gpu-opp-table' does not match
    '^opp-table(-[a-z0-9]+)?$'
    Put the 'gpu' part at the end to match the pattern.

    Fixes: 916a0edc43f0 ("arm64: dts: amlogic: meson-g12: add the Mali OPP table and use
DVFS")
    Reviewed-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Alexander Stein <alexander.stein@mailbox.org>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Link: https://lore.kernel.org/r/20211026182813.900775-2-alexander.stein@mailbox.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 31335cba539bd2a7a6a29ffdceb06289a4a231c1
Author: Jammy Huang <jammy_huang@aspeedtech.com>
Date:   Tue Nov 9 03:12:27 2021 +0000

    media: aspeed: Update signal status immediately to ensure sane hw state

    [ Upstream commit af6d1bde395cac174ee71adcd3fa43f6435c7206 ]

    If res-chg, VE_INTERRUPT_MODE_DETECT_WD irq will be raised. But
    v4l2_input_status won't be updated to no-signal immediately until
    aspeed_video_get_resolution() in aspeed_video_resolution_work().

    During the period of time, aspeed_video_start_frame() could be called
    because it doesn't know signal becomes unstable now. If it goes with
    aspeed_video_init_regs() of aspeed_video_irq_res_change()
    simultaneously, it will mess up hw state.

    To fix this problem, v4l2_input_status is updated to no-signal
    immediately for VE_INTERRUPT_MODE_DETECT_WD irq.

    Fixes: d2b4387f3bdf ("media: platform: Add Aspeed Video Engine driver")
    Signed-off-by: Jammy Huang <jammy_huang@aspeedtech.com>
```

```
    Acked-by: Paul Menzel <pmenzel@molgen.mpg.de>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16641a4deb1acea3b4248d4f219eb645f2c6d7fc
Author: Dongliang Mu <mudongliangabcd@gmail.com>
Date:   Mon Nov 1 09:55:39 2021 +0000

    media: em28xx: fix memory leak in em28xx_init_dev

    [ Upstream commit 22be5a10d0b24eec9e45decd15d7e6112b25f080 ]

    In the em28xx_init_rev, if em28xx_audio_setup fails, this function fails
    to deallocate the media_dev allocated in the em28xx_media_device_init.

    Fix this by adding em28xx_unregister_media_device to free media_dev.

    BTW, this patch is tested in my local syzkaller instance, and it can
    prevent the memory leak from occurring again.

    CC: Pavel Skripkin <paskripkin@gmail.com>
    Fixes: 37ecc7b1278f ("[media] em28xx: add media controller support")
    Signed-off-by: Dongliang Mu <mudongliangabcd@gmail.com>
    Reported-by: syzkaller <syzkaller@googlegroups.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 26ae4318311344cc57551a9917fb1d32d9729130
Author: Jammy Huang <jammy_huang@aspeedtech.com>
Date:   Wed Nov 3 08:23:54 2021 +0000

    media: aspeed: fix mode-detect always time out at 2nd run

    [ Upstream commit 62cea52ad4bead0ae4be2cfe1142eb0aae0e9fbd ]

    aspeed_video_get_resolution() will try to do res-detect again if the
    timing got in last try is invalid. But it will always time out because
    VE_SEQ_CTRL_TRIG_MODE_DET is only cleared after 1st mode-detect.

    To fix the problem, just clear VE_SEQ_CTRL_TRIG_MODE_DET before setting
    it in aspeed_video_enable_mode_detect().

    Fixes: d2b4387f3bdf ("media: platform: Add Aspeed Video Engine driver")
    Signed-off-by: Jammy Huang <jammy_huang@aspeedtech.com>
    Acked-by: Paul Menzel <pmenzel@molgen.mpg.de>
    Reviewed-by: Joel Stanley <joel@jms.id.au>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8203fd3c936375c32736e3e58a5d3e923ae5282
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Tue Oct 12 09:21:50 2021 +0100

    media: atomisp: fix uninitialized bug in gmin_get_pmic_id_and_addr()

    [ Upstream commit cb4d67a998e97365afdf34965b069601da1dae60 ]

    The "power" pointer is not initialized on the else path and that would
    lead to an Oops.

    Link: https://lore.kernel.org/linux-media/20211012082150.GA31086@kili
    Fixes: c30f4cb2d4c7 ("media: atomisp: Refactor PMIC detection to a separate function")
```

```
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b6db0f1f4953747d3a4c167e61b917f9467bcb2
Author: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
Date:   Mon Oct 25 09:06:52 2021 +0100

    media: atomisp: fix enum formats logic

    [ Upstream commit fae46cb0531b45c789e39128f676f2bafa3a7b47 ]

    Changeset 374d62e7aa50 ("media: v4l2-subdev: Verify v4l2_subdev_call() pad config
argument")
    added an extra verification for a pads parameter for enum mbus
    format code.

    Such change broke atomisp, because now the V4L2 core
    refuses to enum MBUS formats if the state is empty.

    So, add .which field in order to select the active formats,
    in order to make it work again.

    While here, improve error messages.

    Fixes: 374d62e7aa50 ("media: v4l2-subdev: Verify v4l2_subdev_call() pad config
argument")
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07ef58acf1093fb01d2a59483bfdf7ff1536421e
Author: Tsuchiya Yuto <kitakar@gmail.com>
Date:   Mon Oct 18 01:23:34 2021 +0900

    media: atomisp: add NULL check for asd obtained from atomisp_video_pipe

    [ Upstream commit c10bcb13462e9cf43111d17f1e08b4bb4d4401b0 ]

    This is almost a BUG report with RFC patch that just avoids kernel
    oopses. Thus, prefixed with [BUG][RFC].

    Here is the kernel log after running `v4l2-compliance -d /dev/video4`
    with this patch applied:

            kern  :err   : [25507.580392] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.592343] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.592995] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.593685] atomisp-isp2 0000:00:03.0: atomisp_g_input(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.593719] atomisp-isp2 0000:00:03.0: atomisp_g_parm(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.593727] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            [omitting 42 same messages]
            kern  :err   : [25507.593976] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594191] atomisp-isp2 0000:00:03.0: atomisp_g_input(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594449] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            [omitting 43 same messages]
            kern  :err   : [25507.594756] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
```

```
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594779] atomisp-isp2 0000:00:03.0: atomisp_g_ctrl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594787] atomisp-isp2 0000:00:03.0: atomisp_s_ctrl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594803] atomisp-isp2 0000:00:03.0:
atomisp_camera_g_ext_ctrls(): asd is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594880] atomisp-isp2 0000:00:03.0: atomisp_enum_fmt_cap():
asd is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.594915] atomisp-isp2 0000:00:03.0: atomisp_g_parm(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.595058] atomisp-isp2 0000:00:03.0: atomisp_try_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.595089] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.595124] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.595221] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.595241] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.601571] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.607496] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.608604] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.611988] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.617420] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.618429] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.618811] atomisp-isp2 0000:00:03.0: atomisp_g_parm(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.622193] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.627355] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.628391] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.631143] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.635813] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.636489] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.636504] atomisp-isp2 0000:00:03.0: atomisp_s_input(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.636516] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.639111] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.646152] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.646831] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.646847] atomisp-isp2 0000:00:03.0: atomisp_s_input(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.650079] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
            kern  :warn  : [25507.657476] isys dma store at addr(0xcd408) val(0)
            kern  :err   : [25507.658741] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.658759] atomisp-isp2 0000:00:03.0: atomisp_s_input(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.658771] atomisp-isp2 0000:00:03.0: atomisp_set_fmt(): asd
is NULL, device is ATOMISP ISP ACC
            kern  :err   : [25507.660959] atomisp-isp2 0000:00:03.0: can't change power
state from D3cold to D0 (config space inaccessible)
```

```
             kern  :warn  : [25507.666665] isys dma store at addr(0xcd408) val(0)
             kern  :err   : [25507.667397] atomisp-isp2 0000:00:03.0: atomisp_queryctl(): asd
 is NULL, device is ATOMISP ISP ACC

    [mchehab: fix coding style]
    Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 1e5817aa63fcfec2673629af077970033cc64016
 Author: Tsuchiya Yuto <kitakar@gmail.com>
 Date:   Mon Oct 18 01:19:47 2021 +0900

    media: atomisp: fix ifdefs in sh_css.c

    [ Upstream commit 5a1b2725558f8a3b4cbf0504f53cffae8e163034 ]

     ## `if (pipe->stream->config.mode == IA_CSS_INPUT_MODE_TPG) {` case

    The intel-aero atomisp has `#if defined(IS_ISP_2400_SYSTEM)` [1]. It is
    to be defined in the following two places [2]:

      - css/hive_isp_css_common/system_global.h
      - css/css_2401_csi2p_system/system_global.h

    and the former file is to be included on ISP2400 devices, too. So, it
    is to be defined for both ISP2400 and ISP2401 devices.

    Because the upstreamed atomisp driver now supports only ISP2400 and
    ISP2401, just remove the ISP version test again. This matches the other
    upstream commits like 3c0538fbad9f ("media: atomisp: get rid of most
    checks for ISP2401 version").

    While here, moved the comment for define GP_ISEL_TPG_MODE to the
    appropriate place.

    [1] https://github.com/intel-aero/linux-
 kernel/blob/a1b673258feb915268377275130c5c5df0eafc82/drivers/media/pci/atomisp/css/sh_css.c#
 L552-L558
    [2] https://github.com/intel-aero/linux-kernel/search?q=IS_ISP_2400_SYSTEM

      ## `isys_stream_descr->polling_mode` case

    This does not exist on the intel-aero atomisp. This is because it is
    based on css version irci_stable_candrpv_0415_20150521_0458.

    On the other hand, the upstreamed atomisp is based on the following css
    version depending on the ISP version using ifdefs:

      - ISP2400: irci_stable_candrpv_0415_20150521_0458
      - ISP2401: irci_master_20150911_0724

    The `isys_stream_descr->polling_mode` usage was added on updating css
    version to irci_master_20150701_0213 [3].

    So, it is not a ISP version specific thing, but css version specific
    thing. Because the upstreamed atomisp driver uses irci_master_20150911_0724
    for ISP2401, re-add the ISP version check for now.

    I say "for now" because ISP2401 should eventually use the same css
    version with ISP2400 (i.e., irci_stable_candrpv_0415_20150521_0458)

    [3] https://raw.githubusercontent.com/intel/ProductionKernelQuilts/cht-m1stable-
 2016_ww31/uefi/cht-m1stable/patches/cam-0439-atomisp2-css2401-and-2401_legacy-
 irci_master_2015070.patch
```

```
         ("atomisp2: css2401 and 2401_legacy-irci_master_20150701_0213")
         Link to Intel's Android kernel patch.

      ## `coord = &me->config.internal_frame_origin_bqs_on_sctbl;` case

     it was added on commit 4f744a573db3 ("media: atomisp: make
     sh_css_sp_init_pipeline() ISP version independent") for ISP2401. Because
     the upstreamed atomisp for the ISP2401 part is based on
     irci_master_20150911_0724, hence the difference.

     Because the upstreamed atomisp driver uses irci_master_20150911_0724
     for ISP2401, revert the test back to `if (IS_ISP2401)`.

     Fixes: 27333dadef57 ("media: atomisp: adjust some code at sh_css that could be broken")
     Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
     Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit bf29ff29e91b09a8791952a4c6f40011a067c060
 Author: Tsuchiya Yuto <kitakar@gmail.com>
 Date:   Mon Oct 18 01:19:45 2021 +0900

     media: atomisp: fix inverted error check for ia_css_mipi_is_source_port_valid()

     [ Upstream commit d21ce8c2f7bf6d737b60c09f86db141b9e8e47f0 ]

     The function ia_css_mipi_is_source_port_valid() returns true if the port
     is valid. So, we can't use the existing err variable as is.

     To fix this issue while reusing that variable, invert the return value
     when assigning it to the variable.

     Fixes: 3c0538fbad9f ("media: atomisp: get rid of most checks for ISP2401 version")
     Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
     Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 3429cc484b72c46172a6c73e49bd8b6abbefc1e9
 Author: Tsuchiya Yuto <kitakar@gmail.com>
 Date:   Mon Oct 18 01:19:44 2021 +0900

     media: atomisp: do not use err var when checking port validity for ISP2400

     [ Upstream commit 9f6b4fa2d2dfbff4b8a57eeb39b1128a6094ee20 ]

     Currently, the `port >= N_CSI_PORTS || err` checks for ISP2400 are always
     evaluated as true because the err variable is set to `-EINVAL` on
     declaration but the variable is never used until the evaluation.

     Looking at the diff of commit 3c0538fbad9f ("media: atomisp: get rid of
     most checks for ISP2401 version"), the `port >= N_CSI_PORTS` check is
     for ISP2400 and the err variable check is for ISP2401. Fix this issue
     by adding ISP version test there accordingly.

     Fixes: 3c0538fbad9f ("media: atomisp: get rid of most checks for ISP2401 version")
     Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
     Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 31449449f1c7fdd42a50ade8da0b8dc254618a93
 Author: Tsuchiya Yuto <kitakar@gmail.com>
 Date:   Mon Oct 18 01:19:43 2021 +0900

     media: atomisp: fix inverted logic in buffers_needed()
```

```
    [ Upstream commit e1921cd14640f0f4d1fad5eb8e448c58a536415d ]

    When config.mode is IA_CSS_INPUT_MODE_BUFFERED_SENSOR, it rather needs
    buffers. Fix it by inverting the return value.

    Fixes: 3c0538fbad9f ("media: atomisp: get rid of most checks for ISP2401 version")
    Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92635e6199516fdd26c534de3431d93c5500e334
Author: Tsuchiya Yuto <kitakar@gmail.com>
Date:   Mon Oct 18 01:19:42 2021 +0900

    media: atomisp: fix punit_ddr_dvfs_enable() argument for mrfld_power up case

    [ Upstream commit 5bfbf65fcca7325e4d89d289b3c286e11220e386 ]

    When comparing with intel-aero atomisp [1], it looks like
    punit_ddr_dvfs_enable() should take `false` as an argument on mrfld_power
    up case.

    Code from the intel-aero kernel [1]:

            int atomisp_mrfld_power_down(struct atomisp_device *isp)
            {
            [...]
                    /*WA:Enable DVFS*/
                    if (IS_CHT)
                            punit_ddr_dvfs_enable(true);

            int atomisp_mrfld_power_up(struct atomisp_device *isp)
            {
            [...]
                    /*WA for PUNIT, if DVFS enabled, ISP timeout observed*/
                    if (IS_CHT)
                            punit_ddr_dvfs_enable(false);

    This patch fixes the inverted argument as per the intel-aero code, as
    well as its comment. While here, fix space issues for comments in
    atomisp_mrfld_power().

    Note that it does not seem to be possible to unify the up/down cases for
    punit_ddr_dvfs_enable(), i.e., we can't do something like the following:

            if (IS_CHT)
                    punit_ddr_dvfs_enable(!enable);

    because according to the intel-aero code [1], the DVFS is disabled
    before "writing 0x0 to ISPSSPM0 bit[1:0]" and the DVFS is enabled after
    "writing 0x3 to ISPSSPM0 bit[1:0]".

    [1] https://github.com/intel-aero/linux-
kernel/blob/a1b673258feb915268377275130c5c5df0eafc82/drivers/media/pci/atomisp/atomisp_drive
r/atomisp_v4l2.c#L431-L514

    Fixes: 0f441fd70b1e ("media: atomisp: simplify the power down/up code")
    Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b5495c5e4b4290cbb02ce40e386d82c6bb60c8c
Author: Tsuchiya Yuto <kitakar@gmail.com>
Date:   Mon Oct 18 01:19:41 2021 +0900
```

```
    media: atomisp: add missing media_device_cleanup() in atomisp_unregister_entities()

    [ Upstream commit ce3015b7212e96db426d0c36f80fd159c91155d1 ]

    After the commit 9832e155f1ed ("[media] media-device: split media
    initialization and registration"), calling media_device_cleanup()
    is needed it seems. However, currently it is missing for the module
    unload path.

    Note that for the probe failure path, it is already added in
    atomisp_register_entities().

    This patch adds the missing call of media_device_cleanup() in
    atomisp_unregister_entities().

    Fixes: a49d25364dfb ("staging/atomisp: Add support for the Intel IPU v2")
    Signed-off-by: Tsuchiya Yuto <kitakar@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c6799722844ecd49b1c5997fe220f5321b1d4f1
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Wed Oct 13 08:53:19 2021 +0100

    media: ipu3-cio2: fix error code in cio2_bridge_connect_sensor()

    [ Upstream commit 85db29d22cc521d9d06de2f5c7832981a55df157 ]

    Return -ENODEV if acpi_get_physical_device_location() fails.  Don't
    return success.

    Fixes: 485aa3df0dff ("media: ipu3-cio2: Parse sensor orientation and rotation")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Daniel Scally <djrscally@gmail.com>
    Reviewed-by: Kieran Bingham <kieran.bingham+renesas@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de0973e0d4f0b0df54405efb17ad1a0e3a6fd217
Author: Dillon Min <dillon.minfei@gmail.com>
Date:   Tue Oct 19 09:43:19 2021 +0100

    media: videobuf2: Fix the size printk format

    [ Upstream commit c9ee220d76775e42f35d634479c978d9350077d3 ]

    Since the type of parameter size is unsigned long,
    it should printk by %lu, instead of %ld, fix it.

    Fixes: 7952be9b6ece ("media: drivers/media/common/videobuf2: rename from videobuf")
    Signed-off-by: Dillon Min <dillon.minfei@gmail.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 872e93c362b58caa7dcc13d00f96f4382dd59638
Author: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
Date:   Mon Oct 25 21:56:28 2021 +0100

    mtd: hyperbus: rpc-if: Check return value of rpcif_sw_init()

    [ Upstream commit 981387ed06b96908223a607f5fba6efa42728fc2 ]

    rpcif_sw_init() can fail so make sure we check the return value
```

```
    of it and on error exit rpcif_hb_probe() callback with error code.

    Fixes: 5de15b610f78 ("mtd: hyperbus: add Renesas RPC-IF driver")
    Signed-off-by: Lad Prabhakar <prabhakar.mahadev-lad.rj@bp.renesas.com>
    Signed-off-by: Vignesh Raghavendra <vigneshr@ti.com>
    Reviewed-by: Biju Das <biju.das.jz@bp.renesas.com>
    Reviewed-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20211025205631.21151-5-prabhakar.mahadev-
lad.rj@bp.renesas.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6035967d49f8980fb75d40c3504c089740ee2371
Author: Quentin Monnet <quentin@isovalent.com>
Date:   Wed Nov 10 11:46:27 2021 +0000

    bpftool: Fix memory leak in prog_dump()

    [ Upstream commit ebbd7f64a3fbe9e0f235e39fc244ee9735e2a52a ]

    Following the extraction of prog_dump() from do_dump(), the struct btf
    allocated in prog_dump() is no longer freed on error; the struct
    bpf_prog_linfo is not freed at all. Make sure we release them before
    exiting the function.

    Fixes: ec2025095cf6 ("bpftool: Match several programs with same tag")
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20211110114632.24537-2-quentin@isovalent.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9047340c8cda9154fc1dbf5897d9b3a68a249c0a
Author: Rameshkumar Sundaram <ramess@codeaurora.org>
Date:   Wed Nov 10 17:10:48 2021 +0200

    ath11k: Send PPDU_STATS_CFG with proper pdev mask to firmware

    [ Upstream commit 16a2c3d5406f95ef6139de52669c60a39443f5f7 ]

    HTT_PPDU_STATS_CFG_PDEV_ID bit mask for target FW PPDU stats request message
    was set as bit 8 to 15. Bit 8 is reserved for soc stats and pdev id starts from
    bit 9. Hence change the bitmask as bit 9 to 15 and fill the proper pdev id in
    the request message.

    In commit 701e48a43e15 ("ath11k: add packet log support for QCA6390"), both
    HTT_PPDU_STATS_CFG_PDEV_ID and pdev_mask were changed, but this pdev_mask
    calculation is not valid for platforms which has multiple pdevs with 1 rxdma
    per pdev, as this is writing same value(i.e. 2) for all pdevs.  Hence fixed it
    to consider pdev_idx as well, to make it compatible for both single and multi
    pd cases.

    Tested on: IPQ8074 hw2.0 AHB WLAN.HK.2.5.0.1-01092-QCAHKSWPL_SILICONZ-1
    Tested on: IPQ6018 hw1.0 WLAN.HK.2.5.0.1-01067-QCAHKSWPL_SILICONZ-1

    Fixes: 701e48a43e15 ("ath11k: add packet log support for QCA6390")

    Co-developed-by: Sathishkumar Muruganandam <murugana@codeaurora.org>
    Signed-off-by: Sathishkumar Muruganandam <murugana@codeaurora.org>
    Signed-off-by: Rameshkumar Sundaram <ramess@codeaurora.org>
    Signed-off-by: Jouni Malinen <jouni@codeaurora.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210721212029.142388-10-jouni@codeaurora.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec0e4b0e42dbd42c120ce422503fd8971568ebc0
```

```
Author: Benjamin Li <benl@squareup.com>
Date:   Wed Nov 3 18:05:48 2021 -0700

    wcn36xx: fix RX BD rate mapping for 5GHz legacy rates

    [ Upstream commit cfdf6b19e750f7de8ae71a26932f63b52e3bf74c ]

    The linear mapping between the BD rate field and the driver's 5GHz
    legacy rates table (wcn_5ghz_rates) does not only apply for the latter
    four rates -- it applies to all eight rates.

    Fixes: 6ea131acea98 ("wcn36xx: Fix warning due to bad rate_idx")
    Signed-off-by: Benjamin Li <benl@squareup.com>
    Tested-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211104010548.1107405-3-benl@squareup.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f32bb8ec62eb82142b9ae4b41636838c954620b3
Author: Benjamin Li <benl@squareup.com>
Date:   Wed Nov 3 18:05:47 2021 -0700

    wcn36xx: populate band before determining rate on RX

    [ Upstream commit c9c5608fafe4dae975c9644c7d14c51ad3b0ed73 ]

    status.band is used in determination of status.rate -- for 5GHz on legacy
    rates there is a linear shift between the BD descriptor's rate field and
    the wcn36xx driver's rate table (wcn_5ghz_rates).

    We have a special clause to populate status.band for hardware scan offload
    frames. However, this block occurs after status.rate is already populated.
    Correctly handle this dependency by moving the band block before the rate
    block.

    This patch addresses kernel warnings & missing scan results for 5GHz APs
    that send their beacons/probe responses at the higher four legacy rates
    (24-54 Mbps), when using hardware scan offload:

      ------------[ cut here ]------------
      WARNING: CPU: 0 PID: 0 at net/mac80211/rx.c:4532 ieee80211_rx_napi+0x744/0x8d8
      Modules linked in: wcn36xx [...]
      CPU: 0 PID: 0 Comm: swapper/0 Tainted: G        W         4.19.107-g73909fa #1
      Hardware name: Square, Inc. T2 (all variants) (DT)
      Call trace:
      dump_backtrace+0x0/0x148
      show_stack+0x14/0x1c
      dump_stack+0xb8/0xf0
      __warn+0x2ac/0x2d8
      warn_slowpath_null+0x44/0x54
      ieee80211_rx_napi+0x744/0x8d8
      ieee80211_tasklet_handler+0xa4/0xe0
      tasklet_action_common+0xe0/0x118
      tasklet_action+0x20/0x28
      __do_softirq+0x108/0x1ec
      irq_exit+0xd4/0xd8
      __handle_domain_irq+0x84/0xbc
      gic_handle_irq+0x4c/0xb8
      el1_irq+0xe8/0x190
      lpm_cpuidle_enter+0x220/0x260
      cpuidle_enter_state+0x114/0x1c0
      cpuidle_enter+0x34/0x48
      do_idle+0x150/0x268
      cpu_startup_entry+0x20/0x24
      rest_init+0xd4/0xe0
```

```
       start_kernel+0x398/0x430
       ---[ end trace ae28cb759352b403 ]---

    Fixes: 8a27ca394782 ("wcn36xx: Correct band/freq reporting on RX")
    Signed-off-by: Benjamin Li <benl@squareup.com>
    Tested-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211104010548.1107405-2-benl@squareup.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd655f6d8a59e615466ac8fc57cb097cc6cde75f
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Fri Nov 5 12:21:52 2021 +0000

    wcn36xx: Put DXE block into reset before freeing memory

    [ Upstream commit ed04ea76e69e7194f7489cebe23a32a68f39218d ]

    When deiniting the DXE hardware we should reset the block to ensure there
    is no spurious DMA write transaction from the downstream WCNSS to upstream
    MSM at a skbuff address we will have released.

    Fixes: 8e84c2582169 ("wcn36xx: mac80211 driver for Qualcomm WCN3660/WCN3680 hardware")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211105122152.1580542-4-bryan.odonoghue@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b1e05391012fdc75e71bf6972e45e5b0e8b879a
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Fri Nov 5 12:21:51 2021 +0000

    wcn36xx: Release DMA channel descriptor allocations

    [ Upstream commit 3652096e5263ad67604b0323f71d133485f410e5 ]

    When unloading the driver we are not releasing the DMA descriptors which we
    previously allocated.

    Fixes: 8e84c2582169 ("wcn36xx: mac80211 driver for Qualcomm WCN3660/WCN3680 hardware")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211105122152.1580542-3-bryan.odonoghue@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c0edbfa92ea97306ee8cf2404770313de01c2466
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Fri Nov 5 12:21:50 2021 +0000

    wcn36xx: Fix DMA channel enable/disable cycle

    [ Upstream commit 89dcb1da611d9b3ff0728502d58372fdaae9ebff ]

    Right now we have a broken sequence where we enable DMA channel interrupts
    which can be left enabled and never disabled if we hit an error path.

    Worse still when we unload the driver, the DMA channel interrupt bits are
    left intact. About the only saving grace here is that we do remember to
    disable the wcnss interrupt when unload the driver.

    Fixes: 8e84c2582169 ("wcn36xx: mac80211 driver for Qualcomm WCN3660/WCN3680 hardware")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211105122152.1580542-2-bryan.odonoghue@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 4d9b9bbe6280d2b37af4758eebabd8cafa0c48b3
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Sun Nov 7 08:55:14 2021 -0800

    libbpf: Free up resources used by inner map definition

    [ Upstream commit 8f7b239ea8cfdc8e64c875ee417fed41431a1f37 ]

    It's not enough to just free(map->inner_map), as inner_map itself can
    have extra memory allocated, like map name.

    Fixes: 646f02ffdd49 ("libbpf: Add BTF-defined map-in-map support")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Reviewed-by: Hengqi Chen <hengqi.chen@gmail.com>
    Link: https://lore.kernel.org/bpf/20211107165521.9240-3-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6309d82228ba56f83b26f377cf9d01e1753b02a
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Fri Nov 5 12:10:55 2021 -0700

    libbpf: Fix non-C89 loop variable declaration in gen_loader.c

    [ Upstream commit b8b5cb55f5d3f03cc1479a3768d68173a10359ad ]

    Fix the `int i` declaration inside the for statement. This is non-C89
    compliant. See [0] for user report breaking BCC build.

      [0] https://github.com/libbpf/libbpf/issues/403

    Fixes: 18f4fccbf314 ("libbpf: Update gen_loader to emit BTF_KIND_FUNC relocations")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/bpf/20211105191055.3324874-1-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04adc4ff421af82fc7b80a8d300748c83de819ae
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Oct 25 17:29:02 2021 +0200

    drm/vc4: hdmi: Enable the scrambler on reconnection

    [ Upstream commit b7551457c5d0b3505b0be247d47919c1ee30506d ]

    If we have a state already and disconnect/reconnect the display, the
    SCDC messages won't be sent again since we didn't go through a disable /
    enable cycle.

    In order to fix this, let's call the vc4_hdmi_enable_scrambling function
    in the detect callback if there is a mode and it needs the scrambler to
    be enabled.

    Fixes: c85695a2016e ("drm/vc4: hdmi: Enable the scrambler")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Link: https://lore.kernel.org/r/20211025152903.1088803-10-maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 658a4485098c01afef48286ae18a98be7a4505bd
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Wed Nov 3 10:32:12 2021 -0700
```

```
    libbpf: Fix section counting logic

    [ Upstream commit 0d6988e16a12ebd41d3e268992211b0ceba44ed7 ]

    e_shnum does include section #0 and as such is exactly the number of ELF
    sections that we need to allocate memory for to use section indices as
    array indices. Fix the off-by-one error.

    This is purely accounting fix, previously we were overallocating one
    too many array items. But no correctness errors otherwise.

    Fixes: 25bbbd7a444b ("libbpf: Remove assumptions about uniqueness of .rodata/.data/.bss
maps")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Acked-by: Yonghong Song <yhs@fb.com>
    Link: https://lore.kernel.org/bpf/20211103173213.1376990-5-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d31d591d44317cf7a20afbab85c70d35d81f4cad
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Thu Oct 28 00:25:29 2021 +0100

    wcn36xx: Indicate beacon not connection loss on MISSED_BEACON_IND

    [ Upstream commit 588b45c88ae130fe373a8c50edaf54735c3f4fe3 ]

    Firmware can trigger a missed beacon indication, this is not the same as a
    lost signal.

    Flag to Linux the missed beacon and let the WiFi stack decide for itself if
    the link is up or down by sending its own probe to determine this.

    We should only be signalling the link is lost when the firmware indicates

    Fixes: 8e84c2582169 ("wcn36xx: mac80211 driver for Qualcomm WCN3660/WCN3680 hardware")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211027232529.657764-1-bryan.odonoghue@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8d34bfcfab32fe38eecbbf68d6c92eafd0ed7131
Author: Benjamin Li <benl@squareup.com>
Date:   Wed Oct 27 10:03:05 2021 -0700

    wcn36xx: ensure pairing of init_scan/finish_scan and start_scan/end_scan

    [ Upstream commit 8f1ba8b0ee2679f0b3d22d2a5c1bc70c436fd872 ]

    An SMD capture from the downstream prima driver on WCN3680B shows the
    following command sequence for connected scans:

    - init_scan_req
        - start_scan_req, channel 1
        - end_scan_req, channel 1
        - start_scan_req, channel 2
        - ...
        - end_scan_req, channel 3
    - finish_scan_req
    - init_scan_req
        - start_scan_req, channel 4
        - ...
        - end_scan_req, channel 6
    - finish_scan_req
    - ...
```

```
        - end_scan_req, channel 165
    - finish_scan_req

    Upstream currently never calls wcn36xx_smd_end_scan, and in some cases[1]
    still sends finish_scan_req twice in a row or before init_scan_req. A
    typical connected scan looks like this:

    - init_scan_req
        - start_scan_req, channel 1
    - finish_scan_req
    - init_scan_req
        - start_scan_req, channel 2
    - ...
        - start_scan_req, channel 165
    - finish_scan_req
    - finish_scan_req

    This patch cleans up scanning so that init/finish and start/end are always
    paired together and correctly nested.

    - init_scan_req
        - start_scan_req, channel 1
        - end_scan_req, channel 1
    - finish_scan_req
    - init_scan_req
        - start_scan_req, channel 2
        - end_scan_req, channel 2
    - ...
        - start_scan_req, channel 165
        - end_scan_req, channel 165
    - finish_scan_req

    Note that upstream will not do batching of 3 active-probe scans before
    returning to the operating channel, and this patch does not change that.
    To match downstream in this aspect, adjust IEEE80211_PROBE_DELAY and/or
    the 125ms max off-channel time in ieee80211_scan_state_decision.

    [1]: commit d195d7aac09b ("wcn36xx: Ensure finish scan is not requested
    before start scan") addressed one case of finish_scan_req being sent
    without a preceding init_scan_req (the case of the operating channel
    coinciding with the first scan channel); two other cases are:
    1) if SW scan is started and aborted immediately, without scanning any
       channels, we send a finish_scan_req without ever sending init_scan_req,
       and
    2) as SW scan logic always returns us to the operating channel before
       calling wcn36xx_sw_scan_complete, finish_scan_req is always sent twice
       at the end of a SW scan

    Fixes: 8e84c2582169 ("wcn36xx: mac80211 driver for Qualcomm WCN3660/WCN3680 hardware")
    Signed-off-by: Benjamin Li <benl@squareup.com>
    Tested-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
    Link: https://lore.kernel.org/r/20211027170306.555535-4-benl@squareup.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b9b45f1726f12c29abeceab3c88cb768b559742
Author: Colin Ian King <colin.king@intel.com>
Date:   Thu Sep 30 11:27:48 2021 +0100

    drm/virtio: fix another potential integer overflow on shift of a int

    [ Upstream commit 74c1bda2f3fa79a93e1c910008649b49b02dc09d ]

    The left shift of unsigned int 32 bit integer constant 1 is evaluated
    using 32 bit arithmetic and then assigned to a signed 64 bit integer.
```

```
    In the case where value is 32 or more this can lead to an overflow
    (value can be in range 0..MAX_CAPSET_ID (63). Fix this by shifting
    the value 1ULL instead.

    Addresses-Coverity: ("Uninitentional integer overflow")
    Fixes: 4fb530e5caf7 ("drm/virtio: implement context init: support init ioctl")
    Signed-off-by: Colin Ian King <colin.king@canonical.com>
    Link: http://patchwork.freedesktop.org/patch/msgid/20210930102748.16922-1-
colin.king@canonical.com
    Signed-off-by: Gerd Hoffmann <kraxel@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4c14c17cf4be58538a86403310159338f7f5477
Author: Colin Ian King <colin.king@intel.com>
Date:   Thu Sep 30 11:19:41 2021 +0100

    drm/virtio: fix potential integer overflow on shift of a int

    [ Upstream commit 8f4502fa284478a5264afa8a5a95511276fa9b80 ]

    The left shift of unsigned int 32 bit integer constant 1 is evaluated
    using 32 bit arithmetic and then assigned to a signed 64 bit integer.
    In the case where i is 32 or more this can lead to an overflow. Fix
    this by shifting the value 1ULL instead.

    Addresses-Coverity: ("Uninitentional integer overflow")
    Fixes: 8d6b006e1f51 ("drm/virtio: implement context init: handle
VIRTGPU_CONTEXT_PARAM_POLL_RINGS_MASK")
    Signed-off-by: Colin Ian King <colin.king@canonical.com>
    Link: http://patchwork.freedesktop.org/patch/msgid/20210930101941.16546-1-
colin.king@canonical.com
    Signed-off-by: Gerd Hoffmann <kraxel@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ac78342d3bab626ebd27cceef606e4502725687
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Mon Oct 25 17:15:28 2021 +0200

    drm/bridge: sn65dsi83: Fix bridge removal

    [ Upstream commit c05f1a4e2c4b8a217b448828c4e59fb47454dc75 ]

    Commit 24417d5b0c00 ("drm/bridge: ti-sn65dsi83: Implement .detach
    callback") moved the unregistration of the bridge DSI device and bridge
    itself to the detach callback.

    While this is correct for the DSI device detach and unregistration, the
    bridge is added in the driver probe, and should thus be removed as part
    of its remove callback.

    Acked-by: Sam Ravnborg <sam@ravnborg.org>
    Reviewed-by: Marek Vasut <marex@denx.de>
    Fixes: 24417d5b0c00 ("drm/bridge: ti-sn65dsi83: Implement .detach callback")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211025151536.1048186-14-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ca720253baad2b137ade7a78a43eda777c2b97f1
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Thu Sep 23 20:50:13 2021 +0200

    drm/vc4: crtc: Make sure the HDMI controller is powered when disabling

    [ Upstream commit bca10db67bdaf15997a5a2a276e7aa9b6eea1393 ]
```

```
    Since commit 875a4d536842 ("drm/vc4: drv: Disable the CRTC at boot
    time"), during the initial setup of the driver we call into the VC4 HDMI
    controller hooks to make sure the controller is properly disabled.

    However, we were never making sure that the device was properly powered
    while doing so. This never resulted in any (reported) issue in practice,
    but since the introduction of commit 4209f03fcb8e ("drm/vc4: hdmi: Warn
    if we access the controller while disabled") we get a loud complaint
    when we do that kind of access.

    Let's make sure we have the HDMI controller properly powered while
    disabling it.

    Fixes: 875a4d536842 ("drm/vc4: drv: Disable the CRTC at boot time")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210923185013.826679-1-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2c3bb9d9c3c8f4b42ae9160e58b3be8f53fae547
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Thu Aug 19 15:59:28 2021 +0200

    drm/vc4: hdmi: Rework the pre_crtc_configure error handling

    [ Upstream commit caa51a4c11f1cadba9bcf61ed9e0105711952ce7 ]

    Since our pre_crtc_configure hook returned void, we didn't implement a
    goto-based error path handling, leading to errors like failing to put
    back the device in pm_runtime in all the error paths, but also failing
    to disable the pixel clock if clk_set_min_rate on the HSM clock fails.

    Move to a goto-based implementation to have an easier consitency.

    Fixes: 4f6e3d66ac52 ("drm/vc4: Add runtime PM support to the HDMI encoder driver")
    Link: https://patchwork.freedesktop.org/patch/msgid/20210819135931.895976-4-
maxime@cerno.tech
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65bf6bfbfb0c0196bbd4a236846d980d93266010
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Thu Aug 19 15:59:27 2021 +0200

    drm/vc4: hdmi: Make sure the controller is powered up during bind

    [ Upstream commit 9c6e4f6ed1d61d5f46946e5c151ceb279eedadb1 ]

    In the bind hook, we actually need the device to have the HSM clock
    running during the final part of the display initialisation where we
    reset the controller and initialise the CEC component.

    Failing to do so will result in a complete, silent, hang of the CPU.

    Fixes: 411efa18e4b0 ("drm/vc4: hdmi: Move the HSM clock enable to runtime_pm")
    Link: https://patchwork.freedesktop.org/patch/msgid/20210819135931.895976-3-
maxime@cerno.tech
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit b3bec971179abc9246840d2cca0532b500a789fa
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Sep 22 14:54:19 2021 +0200

    drm/vc4: hdmi: Make sure the controller is powered in detect

    [ Upstream commit 0f5251339eda7f7eb7bd4467607ae1d01b24e129 ]

    If the HPD GPIO is not available and drm_probe_ddc fails, we end up
    reading the HDMI_HOTPLUG register, but the controller might be powered
    off resulting in a CPU hang. Make sure we have the power domain and the
    HSM clock powered during the detect cycle to prevent the hang from
    happening.

    Fixes: 4f6e3d66ac52 ("drm/vc4: Add runtime PM support to the HDMI encoder driver")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Michael Stapelberg <michael@stapelberg.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210922125419.4125779-6-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1297839fc9dd7ffa33dc8fb39a2a7286a13a39f3
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Sep 22 14:54:18 2021 +0200

    drm/vc4: hdmi: Move the HSM clock enable to runtime_pm

    [ Upstream commit c86b41214362e8e715e1343e16d5d6af0562db05 ]

    In order to access the HDMI controller, we need to make sure the HSM
    clock is enabled. If we were to access it with the clock disabled, the
    CPU would completely hang, resulting in an hard crash.

    Since we have different code path that would require it, let's move that
    clock enable / disable to runtime_pm that will take care of the
    reference counting for us.

    Since we also want to change the HSM clock rate and it's only valid
    while the clock is disabled, we need to move the clk_set_min_rate() call
    on the HSM clock above pm_runtime_get_and_sync().

    Fixes: 4f6e3d66ac52 ("drm/vc4: Add runtime PM support to the HDMI encoder driver")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Michael Stapelberg <michael@stapelberg.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210922125419.4125779-5-
maxime@cerno.tech
    Link: https://lore.kernel.org/linux-arm-kernel/20210924152334.1342630-1-
maxime@cerno.tech/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2a8551369e5f021936680789145 60d6efabc7ad2
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Sep 22 14:54:17 2021 +0200

    drm/vc4: hdmi: Set a default HSM rate

    [ Upstream commit 3e85b81591609bb794bb00cd619b20965b5b38cd ]

    When the firmware doesn't setup the HSM rate (such as when booting
```

```
    without an HDMI cable plugged in), its rate is 0 and thus any register
    access results in a CPU stall, even though HSM is enabled.

    Let's enforce a minimum rate at boot to avoid this issue.

    Fixes: 4f6e3d66ac52 ("drm/vc4: Add runtime PM support to the HDMI encoder driver")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Michael Stapelberg <michael@stapelberg.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210922125419.4125779-4-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 45c127bb0babb66320e8721c0cf8690236ea33e6
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Sep 22 14:54:16 2021 +0200

    clk: bcm-2835: Remove rounding up the dividers

    [ Upstream commit 8ca011ef4af48a7af7b15afd8a4a44039dd04cea ]

    The driver, once it found a divider, tries to round it up by increasing
    the least significant bit of the fractional part by one when the
    round_up argument is set and there's a remainder.

    However, since it increases the divider it will actually reduce the
    clock rate below what we were asking for, leading to issues with
    clk_set_min_rate() that will complain that our rounded clock rate is
    below the minimum of the rate.

    Since the dividers are fairly precise already, let's remove that part so
    that we can have clk_set_min_rate() working.

    This is effectively a revert of 9c95b32ca093 ("clk: bcm2835: add a round
    up ability to the clock divisor").

    Fixes: 9c95b32ca093 ("clk: bcm2835: add a round up ability to the clock divisor")
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Acked-by: Stephen Boyd <sboyd@kernel.org>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org> # boot and basic functionality
    Tested-by: Michael Stapelberg <michael@stapelberg.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210922125419.4125779-3-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a28acf6754276ecdc08f83308dce875df056864e
Author: Maxime Ripard <maxime@cerno.tech>
Date:   Wed Sep 22 14:54:15 2021 +0200

    clk: bcm-2835: Pick the closest clock rate

    [ Upstream commit 5517357a4733d7cf7c17fc79d0530cfa47add372 ]

    The driver currently tries to pick the closest rate that is lower than
    the rate being requested.

    This causes an issue with clk_set_min_rate() since it actively checks
    for the rounded rate to be above the minimum that was just set.

    Let's change the logic a bit to pick the closest rate to the requested
    rate, no matter if it's actually higher or lower.

    Fixes: 6d18b8adbe67 ("clk: bcm2835: Support for clock parent selection")
```

```
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Acked-by: Stephen Boyd <sboyd@kernel.org>
    Reviewed-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Tested-by: Nicolas Saenz Julienne <nsaenz@kernel.org> # boot and basic functionality
    Tested-by: Michael Stapelberg <michael@stapelberg.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210922125419.4125779-2-
maxime@cerno.tech
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 802db7968fadab9a7dbcf3725e36abe30fbca68d
Author: Wang Hai <wanghai38@huawei.com>
Date:   Mon Oct 25 21:10:12 2021 +0800

    Bluetooth: cmtp: fix possible panic when cmtp_init_sockets() fails

    [ Upstream commit 2a7ca7459d905febf519163bd9e3eed894de6bb7 ]

    I got a kernel BUG report when doing fault injection test:

    ------------[ cut here ]------------
    kernel BUG at lib/list_debug.c:45!
    ...
    RIP: 0010:__list_del_entry_valid.cold+0x12/0x4d
    ...
    Call Trace:
     proto_unregister+0x83/0x220
     cmtp_cleanup_sockets+0x37/0x40 [cmtp]
     cmtp_exit+0xe/0x1f [cmtp]
     do_syscall_64+0x35/0xb0
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    If cmtp_init_sockets() in cmtp_init() fails, cmtp_init() still returns
    success. This will cause a kernel bug when accessing uncreated ctmp
    related data when the module exits.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Hai <wanghai38@huawei.com>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad7cb5f6fa5f7ea37208c98a9457dd98025a89ca
Author: Soenke Huster <soenke.huster@eknoes.de>
Date:   Wed Oct 20 10:14:44 2021 +0200

    Bluetooth: virtio_bt: fix memory leak in virtbt_rx_handle()

    [ Upstream commit 1d0688421449718c6c5f46e458a378c9b530ba18 ]

    On the reception of packets with an invalid packet type, the memory of
    the allocated socket buffers is never freed. Add a default case that frees
    these to avoid a memory leak.

    Fixes: afd2daa26c7a ("Bluetooth: Add support for virtio transport driver")
    Signed-off-by: Soenke Huster <soenke.huster@eknoes.de>
    Signed-off-by: Marcel Holtmann <marcel@holtmann.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0881aeaf72e9e824a26eae4ab695a7baf91a91ef
Author: Brian Norris <briannorris@chromium.org>
Date:   Tue Sep 28 14:35:52 2021 -0700

    drm/rockchip: dsi: Disable PLL clock on bind error

    [ Upstream commit 5a614570172e1c9f59035d259dd735acd4f1c01b ]
```

```
    Fix some error handling here noticed in review of other changes.

    Fixes: 2d4f7bdafd70 ("drm/rockchip: dsi: migrate to use dw-mipi-dsi bridge driver")
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reported-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: Chen-Yu Tsai <wenst@chromium.org>
    Tested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210928143413.v3.4.I8bb7a91ecc411d56bc155763f
aa15f289d7fc074@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 479e7dddfec373d73b6cf1af1aae61f8ea0ebb2e
Author: Brian Norris <briannorris@chromium.org>
Date:   Tue Sep 28 14:35:51 2021 -0700

    drm/rockchip: dsi: Fix unbalanced clock on probe error

    [ Upstream commit 251888398753924059f3bb247a44153a2853137f ]

    Our probe() function never enabled this clock, so we shouldn't disable
    it if we fail to probe the bridge.

    Noted by inspection.

    Fixes: 2d4f7bdafd70 ("drm/rockchip: dsi: migrate to use dw-mipi-dsi bridge driver")
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reviewed-by: Chen-Yu Tsai <wenst@chromium.org>
    Tested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210928143413.v3.3.Ie8ceefb51ab6065a1151869b6
fcda41a467d4d2c@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee9cd6d71c8d62b0dfb5e74d2454159264322556
Author: Brian Norris <briannorris@chromium.org>
Date:   Thu Sep 23 17:33:54 2021 -0700

    drm/panel: innolux-p079zca: Delete panel on attach() failure

    [ Upstream commit 32a267e9c057e1636e7afdd20599aa5741a73079 ]

    If we fail to attach (e.g., because 1 of 2 dual-DSI controllers aren't
    ready), we leave a dangling drm_panel reference to freed memory. Clean
    that up on failure.

    This problem exists since the driver's introduction, but is especially
    relevant after refactored for dual-DSI variants.

    Fixes: 14c8f2e9f8ea ("drm/panel: add Innolux P079ZCA panel driver")
    Fixes: 7ad4e4636c54 ("drm/panel: p079zca: Refactor panel driver to support multiple
panels")
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210923173336.2.I9023cf8811a3abf4964ed84eb681
721d8bb489d6@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a98d5747cb61c09b229753308f24062dc14cf85
Author: Brian Norris <briannorris@chromium.org>
Date:   Thu Sep 23 17:33:53 2021 -0700
```

```
    drm/panel: kingdisplay-kd097d04: Delete panel on attach() failure

    [ Upstream commit 5f31dbeae8a88f31c3eb4eb526ab4807c40da241 ]

    If we fail to attach (e.g., because 1 of 2 dual-DSI controllers aren't
    ready), we leave a dangling drm_panel reference to freed memory. Clean
    that up on failure.

    Fixes: 2a994cbed6b2 ("drm/panel: Add Kingdisplay KD097D04 panel driver")
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210923173336.1.Icb4d9dbc1817f4e826361a4f1cea
7461541668f0@changeid
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 83421344496c066ca55ef663ace85003c258ae22
Author: Wang Hai <wanghai38@huawei.com>
Date:   Wed Oct 13 19:41:39 2021 +0800

    drm: fix null-ptr-deref in drm_dev_init_release()

    [ Upstream commit acf20ed020ffa4d6cc8347e8d356509b95df3cbe ]

    I got a null-ptr-deref report:

    [drm:drm_dev_init [drm]] *ERROR* Cannot allocate anonymous inode: -12
    ================================================================
    BUG: KASAN: null-ptr-deref in iput+0x3c/0x4a0
    ...
    Call Trace:
     dump_stack_lvl+0x6c/0x8b
     kasan_report.cold+0x64/0xdb
     __asan_load8+0x69/0x90
     iput+0x3c/0x4a0
     drm_dev_init_release+0x39/0xb0 [drm]
     drm_managed_release+0x158/0x2d0 [drm]
     drm_dev_init+0x3a7/0x4c0 [drm]
     __devm_drm_dev_alloc+0x55/0xd0 [drm]
     mi0283qt_probe+0x8a/0x2b5 [mi0283qt]
     spi_probe+0xeb/0x130
    ...
     entry_SYSCALL_64_after_hwframe+0x44/0xae

    If drm_fs_inode_new() fails in drm_dev_init(), dev->anon_inode will point
    to PTR_ERR(...) instead of NULL. This will result in null-ptr-deref when
    drm_fs_inode_free(dev->anon_inode) is called.

    drm_dev_init()
            drm_fs_inode_new() // fail, dev->anon_inode = PTR_ERR(...)
            drm_managed_release()
                    drm_dev_init_release()
                            drm_fs_inode_free() // access non-existent anon_inode

    Define a temp variable and assign it to dev->anon_inode if the temp
    variable is not PTR_ERR.

    Fixes: 2cbf7fc6718b ("drm: Use drmm_ for drm_dev_init cleanup")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Hai <wanghai38@huawei.com>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211013114139.4042207-1-
wanghai38@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 852226e4aaf10c7afd583c3c7a62431059ffc058
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Wed Oct 13 11:08:25 2021 +0300

    drm/bridge: display-connector: fix an uninitialized pointer in probe()

    [ Upstream commit 189723fbe9aca18d6f7d638c59a40288030932b5 ]

    The "label" pointer is used for debug output.  The code assumes that it
    is either NULL or valid, but it is never set to NULL.  It is either
    valid or uninitialized.

    Fixes: 0c275c30176b ("drm/bridge: Add bridge driver for display connectors")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Sam Ravnborg <sam@ravnborg.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211013080825.GE6010@kili
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f25d782e09b2121b45e6b71a189ce98b61713190
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Thu Oct 7 13:26:25 2021 -0700

    Bluetooth: L2CAP: Fix not initializing sk_peer_pid

    [ Upstream commit f5ff291098f70a70b344df1e388596755c3c8315 ]

    In order to group sockets being connected using L2CAP_MODE_EXT_FLOWCTL
    the pid is used but sk_peer_pid was not being initialized as it is
    currently only done for af_unix.

    Fixes: b48596d1dc25 ("Bluetooth: L2CAP: Add get_peer_pid callback")
    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 ca77b24aac16e633e7ff027fe8e155f151372cec
Author: Tedd Ho-Jeong An <tedd.an@intel.com>
Date:   Wed Oct 6 09:32:28 2021 -0700

    Bluetooth: mgmt: Fix Experimental Feature Changed event

    [ Upstream commit b15bfa4df63529150df9ff0585675f728436e0c1 ]

    This patch fixes the controller index in the Experimental Features
    Changed event for the offload_codec and the quality_report features to
    use the actual hdev index instead of non-controller index(0xffff) so the
    client can receive the event and know which controller the event is for.

    Fixes: ad93315183285 ("Bluetooth: Add offload feature under experimental flag")
    Fixes: ae7d925b5c043 ("Bluetooth: Support the quality report events")
    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 f501b083ffc7b4d829230195d62ffebf70f4db5c
Author: Tedd Ho-Jeong An <tedd.an@intel.com>
Date:   Mon Oct 4 10:01:26 2021 -0700

    Bluetooth: hci_vhci: Fix to set the force_wakeup value

    [ Upstream commit 8b89637dbac2d73d9f3dadf91b4a7dcdb1fc23af ]

    This patch sets the wakeup state of the vhci driver when the
    force_wakeup is updated.
```

```
    Fixes: 60edfad4fd0b6 ("Bluetooth: hci_vhci: Add force_prevent_wake entry")
    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 876a9b33e5e8ffd1043505241b16f4eea181f6e9
Author: xinhui pan <xinhui.pan@amd.com>
Date:   Wed Nov 10 12:31:48 2021 +0800

    drm/ttm: Put BO in its memory manager's lru list

    commit 781050b0a3164934857c300bb0bc291e38c26b6f upstream.

    After we move BO to a new memory region, we should put it to
    the new memory manager's lru list regardless we unlock the resv or not.

    Cc: stable@vger.kernel.org
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: xinhui pan <xinhui.pan@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211110043149.57554-1-
xinhui.pan@amd.com
    Signed-off-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f823374b55a5f306d63de527bb5caf04599b7f79
Author: Brian Norris <briannorris@chromium.org>
Date:   Tue Sep 28 14:35:50 2021 -0700

    drm/rockchip: dsi: Reconfigure hardware on resume()

    commit e584cdc1549932f87a2707b56bc588cfac5d89e0 upstream.

    Since commit 43c2de1002d2 ("drm/rockchip: dsi: move all lane config except
    LCDC mux to bind()"), we perform most HW configuration in the bind()
    function. This configuration may be lost on suspend/resume, so we
    need to call it again. That may lead to errors like this after system
    suspend/resume:

      dw-mipi-dsi-rockchip ff968000.mipi: failed to write command FIFO
      panel-kingdisplay-kd097d04 ff960000.mipi.0: failed write init cmds: -110

    Tested on Acer Chromebook Tab 10 (RK3399 Gru-Scarlet).

    Note that early mailing list versions of this driver borrowed Rockchip's
    downstream/BSP solution, to do HW configuration in mode_set() (which
    *is* called at the appropriate pre-enable() times), but that was
    discarded along the way. I've avoided that still, because mode_set()
    documentation doesn't suggest this kind of purpose as far as I can tell.

    Fixes: 43c2de1002d2 ("drm/rockchip: dsi: move all lane config except LCDC mux to
bind()")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Reviewed-by: Chen-Yu Tsai <wenst@chromium.org>
    Tested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210928143413.v3.2.I4e9d93aadb00b1ffc7d506e31
86a25492bf0b732@changeid
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cbfba9033f143088a9a08912a23ea209680338b0
Author: Brian Norris <briannorris@chromium.org>
Date:   Tue Sep 28 14:35:49 2021 -0700
```

```
    drm/rockchip: dsi: Hold pm-runtime across bind/unbind

    commit 514db871922f103886ad4d221cf406b4fcc5e74a upstream.

    In commit 43c2de1002d2 ("drm/rockchip: dsi: move all lane config except
    LCDC mux to bind()"), we moved most HW configuration to bind(), but we
    didn't move the runtime PM management. Therefore, depending on initial
    boot state, runtime-PM workqueue delays, and other timing factors, we
    may disable our power domain in between the hardware configuration
    (bind()) and when we enable the display. This can cause us to lose
    hardware state and fail to configure our display. For example:

      dw-mipi-dsi-rockchip ff968000.mipi: failed to write command FIFO
      panel-innolux-p079zca ff960000.mipi.0: failed to write command 0

    or:

      dw-mipi-dsi-rockchip ff968000.mipi: failed to write command FIFO
      panel-kingdisplay-kd097d04 ff960000.mipi.0: failed write init cmds: -110

    We should match the runtime PM to the lifetime of the bind()/unbind()
    cycle.

    Tested on Acer Chrometab 10 (RK3399 Gru-Scarlet), with panel drivers
    built either as modules or built-in.

    Side notes: it seems one is more likely to see this problem when the
    panel driver is built into the kernel. I've also seen this problem
    bisect down to commits that simply changed Kconfig dependencies, because
    it changed the order in which driver init functions were compiled into
    the kernel, and therefore the ordering and timing of built-in device
    probe.

    Fixes: 43c2de1002d2 ("drm/rockchip: dsi: move all lane config except LCDC mux to
bind()")
    Link: https://lore.kernel.org/linux-
rockchip/9aedfb528600ecf871885f7293ca4207c84d16c1.camel@gmail.com/
    Reported-by: <aleksandr.o.makarov@gmail.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Brian Norris <briannorris@chromium.org>
    Tested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Chen-Yu Tsai <wenst@chromium.org>
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20210928143413.v3.1.Ic2904d37f30013a7f3d847620
3ad3733c186827e@changeid
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fd7f76e3205408419b381450ba4933fcb85f813f
Author: Gang Li <ligang.bdlg@bytedance.com>
Date:   Fri Jan 14 14:05:23 2022 -0800

    shmem: fix a race between shmem_unused_huge_shrink and shmem_evict_inode

    commit 62c9827cbb996c2c04f615ecd783ce28bcea894b upstream.

    Fix a data race in commit 779750d20b93 ("shmem: split huge pages beyond
    i_size under memory pressure").

    Here are call traces causing race:

       Call Trace 1:
         shmem_unused_huge_shrink+0x3ae/0x410
         ? __list_lru_walk_one.isra.5+0x33/0x160
```

```
          super_cache_scan+0x17c/0x190
          shrink_slab.part.55+0x1ef/0x3f0
          shrink_node+0x10e/0x330
          kswapd+0x380/0x740
          kthread+0xfc/0x130
          ? mem_cgroup_shrink_node+0x170/0x170
          ? kthread_create_on_node+0x70/0x70
          ret_from_fork+0x1f/0x30

         Call Trace 2:
          shmem_evict_inode+0xd8/0x190
          evict+0xbe/0x1c0
          do_unlinkat+0x137/0x330
          do_syscall_64+0x76/0x120
          entry_SYSCALL_64_after_hwframe+0x3d/0xa2

    A simple explanation:

    Image there are 3 items in the local list (@list).  In the first
    traversal, A is not deleted from @list.

      1)    A->B->C
            ^
            |
            pos (leave)

    In the second traversal, B is deleted from @list.  Concurrently, A is
    deleted from @list through shmem_evict_inode() since last reference
    counter of inode is dropped by other thread.  Then the @list is corrupted.

      2)    A->B->C
            ^  ^
            |  |
         evict pos (drop)

    We should make sure the inode is either on the global list or deleted from
    any local list before iput().

    Fixed by moving inodes back to global list before we put them.

    [akpm@linux-foundation.org: coding style fixes]

    Link: https://lkml.kernel.org/r/20211125064502.99983-1-ligang.bdlg@bytedance.com
    Fixes: 779750d20b93 ("shmem: split huge pages beyond i_size under memory pressure")
    Signed-off-by: Gang Li <ligang.bdlg@bytedance.com>
    Reviewed-by: Muchun Song <songmuchun@bytedance.com>
    Acked-by: Kirill A. Shutemov <kirill.shutemov@linux.intel.com>
    Cc: Hugh Dickins <hughd@google.com>
    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 63550aa4935162c73f9c9b25b77ebcfed97efb4e
Author: Wen Gong <quic_wgong@quicinc.com>
Date:   Mon Nov 15 11:29:55 2021 +0200

    ath11k: add string type to search board data in board-2.bin for WCN6855

    commit fc95d10ac41d75c14a81afcc8722333d8b2cf80f upstream.

    Currently ath11k only support string type with bus, chip id and board id
    such as "bus=ahb,qmi-chip-id=1,qmi-board-id=4" for ahb bus chip and
    "bus=pci,qmi-chip-id=0,qmi-board-id=255" for PCIe bus chip in
    board-2.bin. For WCN6855, it is not enough to distinguish all different
```

```
      chips.

      This is to add a new string type which include bus, chip id, board id,
      vendor, device, subsystem-vendor and subsystem-device for WCN6855.

      ath11k will first load board-2.bin and search in it for the board data
      with the above parameters, if matched one board data, then download it
      to firmware, if not matched any one, then ath11k will download the file
      board.bin to firmware.

      Tested-on: WCN6855 hw2.0 PCI WLAN.HSP.1.1-01720.1-QCAHSPSWPL_V1_V2_SILICONZ_LITE-1

      Signed-off-by: Wen Gong <quic_wgong@quicinc.com>
      Signed-off-by: Jouni Malinen <quic_jouni@quicinc.com>
      Signed-off-by: Kalle Valo <kvalo@codeaurora.org>
      Link: https://lore.kernel.org/r/20211111065340.20187-1-quic_wgong@quicinc.com
      Cc: "Limonciello, Mario" <Mario.Limonciello@amd.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 55d524d59bb2dd68aba3db2bf45eba19646d8e8e
  Author: Baoquan He <bhe@redhat.com>
  Date:   Fri Jan 14 14:07:44 2022 -0800

      mm/page_alloc.c: do not warn allocation failure on zone DMA if no managed pages

      commit c4dc63f0032c77464fbd4e7a6afc22fa6913c4a7 upstream.

      In kdump kernel of x86_64, page allocation failure is observed:

       kworker/u2:2: page allocation failure: order:0, mode:0xcc1(GFP_KERNEL|GFP_DMA),
  nodemask=(null),cpuset=/,mems_allowed=0
       CPU: 0 PID: 55 Comm: kworker/u2:2 Not tainted 5.16.0-rc4+ #5
       Hardware name: AMD Dinar/Dinar, BIOS RDN1505B 06/05/2013
       Workqueue: events_unbound async_run_entry_fn
       Call Trace:
        <TASK>
        dump_stack_lvl+0x48/0x5e
        warn_alloc.cold+0x72/0xd6
        __alloc_pages_slowpath.constprop.0+0xc69/0xcd0
        __alloc_pages+0x1df/0x210
        new_slab+0x389/0x4d0
        ___slab_alloc+0x58f/0x770
        __slab_alloc.constprop.0+0x4a/0x80
        kmem_cache_alloc_trace+0x24b/0x2c0
        sr_probe+0x1db/0x620
        ......
        device_add+0x405/0x920
        ......
        __scsi_add_device+0xe5/0x100
        ata_scsi_scan_host+0x97/0x1d0
        async_run_entry_fn+0x30/0x130
        process_one_work+0x1e8/0x3c0
        worker_thread+0x50/0x3b0
        ? rescuer_thread+0x350/0x350
        kthread+0x16b/0x190
        ? set_kthread_struct+0x40/0x40
        ret_from_fork+0x22/0x30
        </TASK>
       Mem-Info:
       ......

      The above failure happened when calling kmalloc() to allocate buffer with
      GFP_DMA.  It requests to allocate slab page from DMA zone while no managed
      pages at all in there.
```

```
      sr_probe()
      --> get_capabilities()
          --> buffer = kmalloc(512, GFP_KERNEL | GFP_DMA);

    Because in the current kernel, dma-kmalloc will be created as long as
    CONFIG_ZONE_DMA is enabled.  However, kdump kernel of x86_64 doesn't have
    managed pages on DMA zone since commit 6f599d84231f ("x86/kdump: Always
    reserve the low 1M when the crashkernel option is specified").  The
    failure can be always reproduced.

    For now, let's mute the warning of allocation failure if requesting pages
    from DMA zone while no managed pages.

    [akpm@linux-foundation.org: fix warning]

    Link: https://lkml.kernel.org/r/20211223094435.248523-4-bhe@redhat.com
    Fixes: 6f599d84231f ("x86/kdump: Always reserve the low 1M when the crashkernel option
is specified")
    Signed-off-by: Baoquan He <bhe@redhat.com>
    Acked-by: John Donnelly  <john.p.donnelly@oracle.com>
    Reviewed-by: Hyeonggon Yoo <42.hyeyoo@gmail.com>
    Cc: Christoph Lameter <cl@linux.com>
    Cc: Pekka Enberg <penberg@kernel.org>
    Cc: David Rientjes <rientjes@google.com>
    Cc: Joonsoo Kim <iamjoonsoo.kim@lge.com>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Cc: Borislav Petkov <bp@alien8.de>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: David Hildenbrand <david@redhat.com>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: Marek Szyprowski <m.szyprowski@samsung.com>
    Cc: Robin Murphy <robin.murphy@arm.com>
    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 8bb3ed1e64a24a573190b52be4af8da0febccea9
Author: Baoquan He <bhe@redhat.com>
Date:   Fri Jan 14 14:07:41 2022 -0800

    dma/pool: create dma atomic pool only if dma zone has managed pages

    commit a674e48c5443d12a8a43c3ac42367aa39505d506 upstream.

    Currently three dma atomic pools are initialized as long as the relevant
    kernel codes are built in.  While in kdump kernel of x86_64, this is not
    right when trying to create atomic_pool_dma, because there's no managed
    pages in DMA zone.  In the case, DMA zone only has low 1M memory
    presented and locked down by memblock allocator.  So no pages are added
    into buddy of DMA zone.  Please check commit f1d4d47c5851 ("x86/setup:
    Always reserve the first 1M of RAM").

    Then in kdump kernel of x86_64, it always prints below failure message:

     DMA: preallocated 128 KiB GFP_KERNEL pool for atomic allocations
     swapper/0: page allocation failure: order:5, mode:0xcc1(GFP_KERNEL|GFP_DMA), nodemask=
(null),cpuset=/,mems_allowed=0
     CPU: 0 PID: 1 Comm: swapper/0 Not tainted 5.13.0-
0.rc5.20210611git929d931f2b40.42.fc35.x86_64 #1
     Hardware name: Dell Inc. PowerEdge R910/0P658H, BIOS 2.12.0 06/04/2018
     Call Trace:
      dump_stack+0x7f/0xa1
      warn_alloc.cold+0x72/0xd6
      __alloc_pages_slowpath.constprop.0+0xf29/0xf50
```

```
       __alloc_pages+0x24d/0x2c0
       alloc_page_interleave+0x13/0xb0
       atomic_pool_expand+0x118/0x210
       __dma_atomic_pool_init+0x45/0x93
       dma_atomic_pool_init+0xdb/0x176
       do_one_initcall+0x67/0x320
       kernel_init_freeable+0x290/0x2dc
       kernel_init+0xa/0x111
       ret_from_fork+0x22/0x30
      Mem-Info:
      ......
      DMA: failed to allocate 128 KiB GFP_KERNEL|GFP_DMA pool for atomic allocation
      DMA: preallocated 128 KiB GFP_KERNEL|GFP_DMA32 pool for atomic allocations

    Here, let's check if DMA zone has managed pages, then create
    atomic_pool_dma if yes.  Otherwise just skip it.

    Link: https://lkml.kernel.org/r/20211223094435.248523-3-bhe@redhat.com
    Fixes: 6f599d84231f ("x86/kdump: Always reserve the low 1M when the crashkernel option
 is specified")
    Signed-off-by: Baoquan He <bhe@redhat.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Acked-by: John Donnelly  <john.p.donnelly@oracle.com>
    Reviewed-by: David Hildenbrand <david@redhat.com>
    Cc: Marek Szyprowski <m.szyprowski@samsung.com>
    Cc: Robin Murphy <robin.murphy@arm.com>
    Cc: Borislav Petkov <bp@alien8.de>
    Cc: Christoph Lameter <cl@linux.com>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: David Rientjes <rientjes@google.com>
    Cc: Hyeonggon Yoo <42.hyeyoo@gmail.com>
    Cc: Joonsoo Kim <iamjoonsoo.kim@lge.com>
    Cc: Pekka Enberg <penberg@kernel.org>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    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 e782d7a1b03be103a56fefb8d7be906c7f527f23
 Author: Baoquan He <bhe@redhat.com>
 Date:   Fri Jan 14 14:07:37 2022 -0800

    mm_zone: add function to check if managed dma zone exists

    commit 62b3107073646e0946bd97ff926832bafb846d17 upstream.

    Patch series "Handle warning of allocation failure on DMA zone w/o
    managed pages", v4.

    **Problem observed:
    On x86_64, when crash is triggered and entering into kdump kernel, page
    allocation failure can always be seen.

     ---------------------------------
     DMA: preallocated 128 KiB GFP_KERNEL pool for atomic allocations
     swapper/0: page allocation failure: order:5, mode:0xcc1(GFP_KERNEL|GFP_DMA), nodemask=
 (null),cpuset=/,mems_allowed=0
     CPU: 0 PID: 1 Comm: swapper/0
     Call Trace:
      dump_stack+0x7f/0xa1
      warn_alloc.cold+0x72/0xd6
      ......
      __alloc_pages+0x24d/0x2c0
      ......
```

```
       dma_atomic_pool_init+0xdb/0x176
       do_one_initcall+0x67/0x320
       ? rcu_read_lock_sched_held+0x3f/0x80
       kernel_init_freeable+0x290/0x2dc
       ? rest_init+0x24f/0x24f
       kernel_init+0xa/0x111
       ret_from_fork+0x22/0x30
      Mem-Info:
      -----------------------------------
```

```
    ***Root cause:
    In the current kernel, it assumes that DMA zone must have managed pages
    and try to request pages if CONFIG_ZONE_DMA is enabled. While this is not
    always true. E.g in kdump kernel of x86_64, only low 1M is presented and
    locked down at very early stage of boot, so that this low 1M won't be
    added into buddy allocator to become managed pages of DMA zone. This
    exception will always cause page allocation failure if page is requested
    from DMA zone.

    ***Investigation:
    This failure happens since below commit merged into linus's tree.
      1a6a9044b967 x86/setup: Remove CONFIG_X86_RESERVE_LOW and reservelow= options
      23721c8e92f7 x86/crash: Remove crash_reserve_low_1M()
      f1d4d47c5851 x86/setup: Always reserve the first 1M of RAM
      7c321eb2b843 x86/kdump: Remove the backup region handling
      6f599d84231f x86/kdump: Always reserve the low 1M when the crashkernel option is
 specified

    Before them, on x86_64, the low 640K area will be reused by kdump kernel.
    So in kdump kernel, the content of low 640K area is copied into a backup
    region for dumping before jumping into kdump. Then except of those firmware
    reserved region in [0, 640K], the left area will be added into buddy
    allocator to become available managed pages of DMA zone.

    However, after above commits applied, in kdump kernel of x86_64, the low
    1M is reserved by memblock, but not released to buddy allocator. So any
    later page allocation requested from DMA zone will fail.

    At the beginning, if crashkernel is reserved, the low 1M need be locked
    down because AMD SME encrypts memory making the old backup region
    mechanims impossible when switching into kdump kernel.

    Later, it was also observed that there are BIOSes corrupting memory
    under 1M. To solve this, in commit f1d4d47c5851, the entire region of
    low 1M is always reserved after the real mode trampoline is allocated.

    Besides, recently, Intel engineer mentioned their TDX (Trusted domain
    extensions) which is under development in kernel also needs to lock down
    the low 1M. So we can't simply revert above commits to fix the page allocation
    failure from DMA zone as someone suggested.

    ***Solution:
    Currently, only DMA atomic pool and dma-kmalloc will initialize and
    request page allocation with GFP_DMA during bootup.

    So only initializ DMA atomic pool when DMA zone has available managed
    pages, otherwise just skip the initialization.

    For dma-kmalloc(), for the time being, let's mute the warning of
    allocation failure if requesting pages from DMA zone while no manged
    pages.  Meanwhile, change code to use dma_alloc_xx/dma_map_xx API to
    replace kmalloc(GFP_DMA), or do not use GFP_DMA when calling kmalloc() if
    not necessary.  Christoph is posting patches to fix those under
    drivers/scsi/.  Finally, we can remove the need of dma-kmalloc() as people
    suggested.
```

```
    This patch (of 3):

    In some places of the current kernel, it assumes that dma zone must have
    managed pages if CONFIG_ZONE_DMA is enabled.  While this is not always
    true.  E.g in kdump kernel of x86_64, only low 1M is presented and locked
    down at very early stage of boot, so that there's no managed pages at all
    in DMA zone.  This exception will always cause page allocation failure if
    page is requested from DMA zone.

    Here add function has_managed_dma() and the relevant helper functions to
    check if there's DMA zone with managed pages.  It will be used in later
    patches.

    Link: https://lkml.kernel.org/r/20211223094435.248523-1-bhe@redhat.com
    Link: https://lkml.kernel.org/r/20211223094435.248523-2-bhe@redhat.com
    Fixes: 6f599d84231f ("x86/kdump: Always reserve the low 1M when the crashkernel option
is specified")
    Signed-off-by: Baoquan He <bhe@redhat.com>
    Reviewed-by: David Hildenbrand <david@redhat.com>
    Acked-by: John Donnelly  <john.p.donnelly@oracle.com>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Christoph Lameter <cl@linux.com>
    Cc: Hyeonggon Yoo <42.hyeyoo@gmail.com>
    Cc: Pekka Enberg <penberg@kernel.org>
    Cc: David Rientjes <rientjes@google.com>
    Cc: Joonsoo Kim <iamjoonsoo.kim@lge.com>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: Borislav Petkov <bp@alien8.de>
    Cc: Marek Szyprowski <m.szyprowski@samsung.com>
    Cc: Robin Murphy <robin.murphy@arm.com>
    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 febf7b7b68666b08fc9a9c86b30624bddf91a3f8
Author: Yifeng Li <tomli@tomli.me>
Date:   Thu Dec 2 06:35:21 2021 +0000

    PCI: Add function 1 DMA alias quirk for Marvell 88SE9125 SATA controller

    commit e445375882883f69018aa669b67cbb37ec873406 upstream.

    Like other SATA controller chips in the Marvell 88SE91xx series, the
    Marvell 88SE9125 has the same DMA requester ID hardware bug that prevents
    it from working under IOMMU.  Add it to the list of devices that need the
    quirk.

    Without this patch, device initialization fails with DMA errors:

      ata8: softreset failed (1st FIS failed)
      DMAR: DRHD: handling fault status reg 2
      DMAR: [DMA Write NO_PASID] Request device [03:00.1] fault addr 0xfffc0000 [fault
reason 0x02] Present bit in context entry is clear
      DMAR: DRHD: handling fault status reg 2
      DMAR: [DMA Read NO_PASID] Request device [03:00.1] fault addr 0xfffc0000 [fault reason
0x02] Present bit in context entry is clear

    After applying the patch, the controller can be successfully initialized:

      ata8: SATA link up 1.5 Gbps (SStatus 113 SControl 330)
      ata8.00: ATAPI: PIONEER BD-RW   BDR-207M, 1.21, max UDMA/100
      ata8.00: configured for UDMA/100
```

```
       scsi 7:0:0:0: CD-ROM            PIONEER  BD-RW   BDR-207M 1.21 PQ: 0 ANSI: 5

    Link: https://lore.kernel.org/r/YahpKVR+McJVDdkD@work
    Reported-by: Sam Bingner <sam@bingner.com>
    Tested-by: Sam Bingner <sam@bingner.com>
    Tested-by: Yifeng Li <tomli@tomli.me>
    Signed-off-by: Yifeng Li <tomli@tomli.me>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Krzysztof Wilczyński <kw@linux.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6596f4ee904ffff0a70fbab1b918acc65d842d85
Author: Thomas Hellström <thomas.hellstrom@linux.intel.com>
Date:   Mon Nov 29 16:27:27 2021 +0100

    dma_fence_array: Fix PENDING_ERROR leak in dma_fence_array_signaled()

    commit 95d35838880fb040ccb9fe4a48816bd0c8b62df5 upstream.

    If a dma_fence_array is reported signaled by a call to
    dma_fence_is_signaled(), it may leak the PENDING_ERROR status.

    Fix this by clearing the PENDING_ERROR status if we return true in
    dma_fence_array_signaled().

    v2:
    - Update Cc list, and add R-b.

    Fixes: 1f70b8b812f3 ("dma-fence: Propagate errors to dma-fence-array container")
    Cc: Chris Wilson <chris@chris-wilson.co.uk>
    Cc: Sumit Semwal <sumit.semwal@linaro.org>
    Cc: Gustavo Padovan <gustavo@padovan.org>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: "Christian König" <christian.koenig@amd.com>
    Cc: linux-media@vger.kernel.org
    Cc: dri-devel@lists.freedesktop.org
    Cc: linaro-mm-sig@lists.linaro.org
    Cc: <stable@vger.kernel.org> # v5.4+
    Signed-off-by: Thomas Hellström <thomas.hellstrom@linux.intel.com>
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20211129152727.448908-1-
thomas.hellstrom@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d73f2b69ee644aeb459a7b90bca18426bd00b146
Author: Peng Hao <flyingpenghao@gmail.com>
Date:   Wed Dec 22 09:12:25 2021 +0800

    virtio/virtio_mem: handle a possible NULL as a memcpy parameter

    commit cf4a4493ff70874f8af26d75d4346c591c298e89 upstream.

    There is a check for vm->sbm.sb_states before, and it should check
    it here as well.

    Signed-off-by: Peng Hao <flyingpeng@tencent.com>
    Link: https://lore.kernel.org/r/20211222011225.40573-1-flyingpeng@tencent.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Fixes: 5f1f79bbc9e2 ("virtio-mem: Paravirtualized memory hotplug")
    Cc: stable@vger.kernel.org # v5.8+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c1acb2919c76169e7179b4e4df5c4b8e1f8489f5
Author: Hao Xu <haoxu@linux.alibaba.com>
```

```
Date:   Thu Nov 25 17:21:02 2021 +0800

    io_uring: fix no lock protection for ctx->cq_extra

    commit e302f1046f4c209291b07ff7bc4d15ca26891f16 upstream.

    ctx->cq_extra should be protected by completion lock so that the
    req_need_defer() does the right check.

    Cc: stable@vger.kernel.org
    Signed-off-by: Hao Xu <haoxu@linux.alibaba.com>
    Link: https://lore.kernel.org/r/20211125092103.224502-2-haoxu@linux.alibaba.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8f08fa2b7584e67f3a94afd2c90a942e2cd99518
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Sat Dec 4 17:58:49 2021 +0300

    drm/tegra: Add back arm_iommu_detach_device()

    commit d210919dbdc8a82c676cc3e3c370b1802be63124 upstream.

    DMA buffers of 2D/3D engines aren't mapped properly when
    CONFIG_ARM_DMA_USE_IOMMU=y. The memory management code of Tegra DRM driver
    has a longstanding overhaul overdue and it's not obvious where the problem
    is in this case. Hence let's add back the old workaround which we already
    had sometime before. It explicitly detaches DRM devices from the offending
    implicit IOMMU domain. This fixes a completely broken 2d/3d drivers in
    case of ARM32 multiplatform kernel config.

    Cc: stable@vger.kernel.org
    Fixes: fa6661b7aa0b ("drm/tegra: Optionally attach clients to the IOMMU")
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit deaca7c1a3d98b24a5162ca26f80542c7fc1fc81
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Sat Dec 4 17:58:48 2021 +0300

    gpu: host1x: Add back arm_iommu_detach_device()

    commit d5185965c3b59073c4520bad7dd2adf725b9abba upstream.

    Host1x DMA buffer isn't mapped properly when CONFIG_ARM_DMA_USE_IOMMU=y.
    The memory management code of Host1x driver has a longstanding overhaul
    overdue and it's not obvious where the problem is in this case. Hence
    let's add back the old workaround which we already had sometime before.
    It explicitly detaches Host1x device from the offending implicit IOMMU
    domain. This fixes a completely broken Host1x DMA in case of ARM32
    multiplatform kernel config.

    Cc: stable@vger.kernel.org
    Fixes: af1cbfb9bf0f ("gpu: host1x: Support DMA mapping of buffers")
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4c8806184228ccbda0d7e765908c471129962046
Author: Yunfei Wang <yf.wang@mediatek.com>
Date:   Tue Dec 7 19:33:15 2021 +0800

    iommu/io-pgtable-arm-v7s: Add error handle for page table allocation failure
```

```
    commit a556cfe4cabc6d79cbb7733f118bbb420b376fe6 upstream.

    In __arm_v7s_alloc_table function:
    iommu call kmem_cache_alloc to allocate page table, this function
    allocate memory may fail, when kmem_cache_alloc fails to allocate
    table, call virt_to_phys will be abnomal and return unexpected phys
    and goto out_free, then call kmem_cache_free to release table will
    trigger KE, __get_free_pages and free_pages have similar problem,
    so add error handle for page table allocation failure.

    Fixes: 29859aeb8a6e ("iommu/io-pgtable-arm-v7s: Abort allocation when table address
overflows the PTE")
    Signed-off-by: Yunfei Wang <yf.wang@mediatek.com>
    Cc: <stable@vger.kernel.org> # 5.10.*
    Acked-by: Robin Murphy <robin.murphy@arm.com>
    Link: https://lore.kernel.org/r/20211207113315.29109-1-yf.wang@mediatek.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5f61db6921e81de57b4453fa0075fb658d9de85a
Author: Hari Prasath <Hari.PrasathGE@microchip.com>
Date:   Wed Dec 8 12:05:53 2021 +0530

    ARM: dts: at91: update alternate function of signal PD20

    commit 12f332d2dd3187472f595b678246adb10d886bd0 upstream.

    The alternate function of PD20 is 4 as per the datasheet of
    sama7g5 and not 5 as defined earlier.

    Signed-off-by: Hari Prasath <Hari.PrasathGE@microchip.com>
    Fixes: 7540629e2fc7 ("ARM: dts: at91: add sama7g5 SoC DT and sama7g5-ek")
    Cc: <stable@vger.kernel.org> # v5.15+
    Signed-off-by: Nicolas Ferre <nicolas.ferre@microchip.com>
    Link: https://lore.kernel.org/r/20211208063553.19807-1-Hari.PrasathGE@microchip.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5e8f710b9af827966c059d0aed663bb1ae37dd45
Author: D Scott Phillips <scott@os.amperecomputing.com>
Date:   Mon Dec 20 15:41:14 2021 -0800

    arm64: errata: Fix exec handling in erratum 1418040 workaround

    commit 38e0257e0e6f4fef2aa2966b089b56a8b1cfb75c upstream.

    The erratum 1418040 workaround enables CNTVCT_EL1 access trapping in EL0
    when executing compat threads. The workaround is applied when switching
    between tasks, but the need for the workaround could also change at an
    exec(), when a non-compat task execs a compat binary or vice versa. Apply
    the workaround in arch_setup_new_exec().

    This leaves a small window of time between SET_PERSONALITY and
    arch_setup_new_exec where preemption could occur and confuse the old
    workaround logic that compares TIF_32BIT between prev and next. Instead, we
    can just read cntkctl to make sure it's in the state that the next task
    needs. I measured cntkctl read time to be about the same as a mov from a
    general-purpose register on N1. Update the workaround logic to examine the
    current value of cntkctl instead of the previous task's compat state.

    Fixes: d49f7d7376d0 ("arm64: Move handling of erratum 1418040 into C code")
    Cc: <stable@vger.kernel.org> # 5.9.x
    Signed-off-by: D Scott Phillips <scott@os.amperecomputing.com>
    Reviewed-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20211220234114.3926-1-scott@os.amperecomputing.com
    Signed-off-by: Catalin Marinas <catalin.marinas@arm.com>
```

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

commit 7d0a23de5f41810827891c844a6ab7653bda4028
Author: Sumeet Pawnikar <sumeet.r.pawnikar@intel.com>
Date:   Thu Dec 23 15:12:36 2021 +0530

    thermal/drivers/int340x: Fix RFIM mailbox write commands

    commit 2685c77b80a80c57e2a25a726b82fb31e6e212ab upstream.

    The existing mail mechanism only supports writing of workload types.

    However, mailbox command for RFIM (cmd = 0x08) also requires write
    operation which is ignored. This results in failing to store RFI
    restriction.

    Fixint this requires enhancing mailbox writes for non workload
    commands too, so remove the check for MBOX_CMD_WORKLOAD_TYPE_WRITE
    in mailbox write to allow this other write commands to be supoorted.

    At the same time, however, we have to make sure that there is no
    impact on read commands, by avoiding to write anything into the
    mailbox data register.

    To properly implement that, add two separate functions for mbox read
    and write commands for the processor thermal workload command type.
    This helps to distinguish the read and write workload command types
    from each other while sending mbox commands.

    Fixes: 5d6fbc96bd36 ("thermal/drivers/int340x: processor_thermal: Export additional
attributes")
    Signed-off-by: Sumeet Pawnikar <sumeet.r.pawnikar@intel.com>
    Cc: 5.14+ <stable@vger.kernel.org> # 5.14+
    Acked-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    [ rjw: Changelog edits ]
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 31fec36490b2f28387d7339f4517b8331e18112c
Author: Dan Williams <dan.j.williams@intel.com>
Date:   Thu Nov 11 10:19:05 2021 -0800

    cxl/pmem: Fix module reload vs workqueue state

    commit 53989fad1286e652ea3655ae3367ba698da8d2ff upstream.

    A test of the form:

        while true; do modprobe -r cxl_pmem; modprobe cxl_pmem; done

    May lead to a crash signature of the form:

        BUG: unable to handle page fault for address: ffffffffc0660030
        #PF: supervisor instruction fetch in kernel mode
        #PF: error_code(0x0010) - not-present page
        [..]
        Workqueue: cxl_pmem 0xffffffffc0660030
        RIP: 0010:0xffffffffc0660030
        Code: Unable to access opcode bytes at RIP 0xffffffffc0660006.
        [..]
        Call Trace:
         ? process_one_work+0x4ec/0x9c0
         ? pwq_dec_nr_in_flight+0x100/0x100
         ? rwlock_bug.part.0+0x60/0x60
         ? worker_thread+0x2eb/0x700
```

```
    In that report the 0xffffffffc0660030 address corresponds to the former
    function address of cxl_nvb_update_state() from a previous load of the
    module, not the current address. Fix that by arranging for ->state_work
    in the 'struct cxl_nvdimm_bridge' object to be reinitialized on cxl_pmem
    module reload.

    Details:

    Recall that CXL subsystem wants to link a CXL memory expander device to
    an NVDIMM sub-hierarchy when both a persistent memory range has been
    registered by the CXL platform driver (cxl_acpi) *and* when that CXL
    memory expander has published persistent memory capacity (Get Partition
    Info). To this end the cxl_nvdimm_bridge driver arranges to rescan the
    CXL bus when either of those conditions change. The helper
    bus_rescan_devices() can not be called underneath the device_lock() for
    any device on that bus, so the cxl_nvdimm_bridge driver uses a workqueue
    for the rescan.

    Typically a driver allocates driver data to hold a 'struct work_struct'
    for a driven device, but for a workqueue that may run after ->remove()
    returns, driver data will have been freed. The 'struct
    cxl_nvdimm_bridge' object holds the state and work_struct directly.
    Unfortunately it was only arranging for that infrastructure to be
    initialized once per device creation rather than the necessary once per
    workqueue (cxl_pmem_wq) creation.

    Introduce is_cxl_nvdimm_bridge() and cxl_nvdimm_bridge_reset() in
    support of invalidating stale references to a recently destroyed
    cxl_pmem_wq.

    Cc: <stable@vger.kernel.org>
    Fixes: 8fdcb1704f61 ("cxl/pmem: Add initial infrastructure for pmem support")
    Reported-by: Vishal Verma <vishal.l.verma@intel.com>
    Tested-by: Vishal Verma <vishal.l.verma@intel.com>
    Link: https://lore.kernel.org/r/163665474585.3505991.8397182770066720755.stgit@dwillia2-
desk3.amr.corp.intel.com
    Signed-off-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c676c00cea441a4e72df5e36ca7397d95421d070
Author: Dan Williams <dan.j.williams@intel.com>
Date:   Fri Oct 29 12:55:47 2021 -0700

    cxl/pmem: Fix reference counting for delayed work

    commit 08b9e0ab8af48895337192e683de44ab1e1b7427 upstream.

    There is a potential race between queue_work() returning and the
    queued-work running that could result in put_device() running before
    get_device(). Introduce the cxl_nvdimm_bridge_state_work() helper that
    takes the reference unconditionally, but drops it if no new work was
    queued, to keep the references balanced.

    Fixes: 8fdcb1704f61 ("cxl/pmem: Add initial infrastructure for pmem support")
    Cc: <stable@vger.kernel.org>
    Reviewed-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Reviewed-by: Ben Widawsky <ben.widawsky@intel.com>
    Link: https://lore.kernel.org/r/163553734757.2509761.3305231863616785470.stgit@dwillia2-
desk3.amr.corp.intel.com
    Signed-off-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8258ffea97b84cd5aa8fbf5a071214197a8bec82
Author: Manivannan Sadhasivam <mani@kernel.org>
```

```
Date:   Thu Dec 16 13:42:24 2021 +0530

    bus: mhi: core: Fix race while handling SYS_ERR at power up

    commit d651ce8e917fa1bf6cfab8dca74c512edffc35d3 upstream.

    During SYS_ERR condition, as a response to the MHI_RESET from host, some
    devices tend to issue BHI interrupt without clearing the SYS_ERR state in
    the device. This creates a race condition and causes a failure in booting
    up the device.

    The issue is seen on the Sierra Wireless EM9191 modem during SYS_ERR
    handling in mhi_async_power_up(). Once the host detects that the device
    is in SYS_ERR state, it issues MHI_RESET and waits for the device to
    process the reset request. During this time, the device triggers the BHI
    interrupt to the host without clearing SYS_ERR condition. So the host
    starts handling the SYS_ERR condition again.

    To fix this issue, let's register the IRQ handler only after handling the
    SYS_ERR check to avoid getting spurious IRQs from the device.

    Fixes: e18d4e9fa79b ("bus: mhi: core: Handle syserr during power_up")
    Cc: stable@vger.kernel.org
    Reported-by: Aleksander Morgado <aleksander@aleksander.es>
    Tested-by: Aleksander Morgado <aleksander@aleksander.es>
    Tested-by: Thomas Perrot <thomas.perrot@bootlin.com>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20211216081227.237749-8-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bdbbfc0d8b78a6ffae14b00c5cc4605d028f4390
Author: Bhaumik Bhatt <quic_bbhatt@quicinc.com>
Date:   Thu Dec 16 13:42:23 2021 +0530

    bus: mhi: core: Fix reading wake_capable channel configuration

    commit 42c4668f7efe1485dfc382517b412c0c6ab102b8 upstream.

    The 'wake-capable' entry in channel configuration is not set when
    parsing the configuration specified by the controller driver. Add
    the missing entry to ensure channel is correctly specified as a
    'wake-capable' channel.

    Link: https://lore.kernel.org/r/1638320491-13382-1-git-send-email-
quic_bbhatt@quicinc.com
    Fixes: 0cbf260820fa ("bus: mhi: core: Add support for registering MHI controllers")
    Cc: stable@vger.kernel.org
    Reviewed-by: Manivannan Sadhasivam <mani@kernel.org>
    Signed-off-by: Bhaumik Bhatt <quic_bbhatt@quicinc.com>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20211216081227.237749-7-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit abee0f615e8c9c75543f290dafad8cf9f34e8128
Author: Loic Poulain <loic.poulain@linaro.org>
Date:   Thu Dec 16 13:42:19 2021 +0530

    bus: mhi: pci_generic: Graceful shutdown on freeze

    commit f77097ec8c0141a4b5cf3722a246be0cb5677e29 upstream.

    There is no reason for shutting down MHI ungracefully on freeze,
    this causes the MHI host stack & device stack to not be aligned
    anymore since the proper MHI reset sequence is not performed for
    ungraceful shutdown.
```

```
    Link: https://lore.kernel.org/r/1635268180-13699-1-git-send-email-
loic.poulain@linaro.org
    Fixes: 5f0c2ee1fe8d ("bus: mhi: pci-generic: Fix hibernation")
    Cc: stable@vger.kernel.org
    Suggested-by: Bhaumik Bhatt <bbhatt@codeaurora.org>
    Reviewed-by: Bhaumik Bhatt <bbhatt@codeaurora.org>
    Reviewed-by: Hemant Kumar <hemantk@codeaurora.org>
    Reviewed-by: Manivannan Sadhasivam <mani@kernel.org>
    Signed-off-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20211216081227.237749-3-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 89421c976d5c34e4420255e5950bd263ebe2cfa8
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Fri Oct 8 18:58:40 2021 +0200

    lkdtm: Fix content of section containing lkdtm_rodata_do_nothing()

    commit bc93a22a19eb2b68a16ecf04cdf4b2ed65aaf398 upstream.

    On a kernel without CONFIG_STRICT_KERNEL_RWX, running EXEC_RODATA
    test leads to "Illegal instruction" failure.

    Looking at the content of rodata_objcopy.o, we see that the
    function content zeroes only:

            Disassembly of section .rodata:

            0000000000000000 <.lkdtm_rodata_do_nothing>:
               0:   00 00 00 00     .long 0x0

    Add the contents flag in order to keep the content of the section
    while renaming it.

            Disassembly of section .rodata:

            0000000000000000 <.lkdtm_rodata_do_nothing>:
               0:   4e 80 00 20     blr

    Fixes: e9e08a07385e ("lkdtm: support llvm-objcopy")
    Cc: stable@vger.kernel.org
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Arnd Bergmann <arnd@arndb.de>
    Cc: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link:
https://lore.kernel.org/r/8900731fbc05fb8b0de18af7133a8fc07c3c53a1.1633712176.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2c5c948451d57596640b20887b06b523d376eb24
Author: Jonathan Cameron <Jonathan.Cameron@huawei.com>
Date:   Sun Oct 17 18:22:09 2021 +0100

    iio: trigger: Fix a scheduling whilst atomic issue seen on tsc2046

    commit 9020ef659885f2622cfb386cc229b6d618362895 upstream.

    IIO triggers are software IRQ chips that split an incoming IRQ into
```

```
    separate IRQs routed to all devices using the trigger.
    When all consumers are done then a trigger callback reenable() is
    called.  There are a few circumstances under which this can happen
    in atomic context.

    1) A single user of the trigger that calls the iio_trigger_done()
    function from interrupt context.
    2) A race between disconnecting the last device from a trigger and
    the trigger itself sucessfully being disabled.

    To avoid a resulting scheduling whilst atomic, close this second corner
    by using schedule_work() to ensure the reenable is not done in atomic
    context.

    Note that drivers must be careful to manage the interaction of
    set_state() and reenable() callbacks to ensure appropriate reference
    counting if they are relying on the same hardware controls.

    Deliberately taking this the slow path rather than via a fixes tree
    because the error has hard to hit and I would like it to soak for a while
    before hitting a release kernel.

    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Cc: Pengutronix Kernel Team <kernel@pengutronix.de>
    Cc: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Tested-by: Oleksij Rempel <o.rempel@pengutronix.de>
    Cc: <Stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20211017172209.112387-1-jic23@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e5cf4b452f8e0647b1a6e260a2db18be698b53c0
Author: Jonathan Cameron <Jonathan.Cameron@huawei.com>
Date:   Sun Dec 5 17:27:28 2021 +0000

    iio: adc: ti-adc081c: Partial revert of removal of ACPI IDs

    commit c9791a94384af07592d29504004d2255dbaf8663 upstream.

    Unfortuanately a non standards compliant ACPI ID is known to be
    in the wild on some AAEON boards.

    Partly revert the removal of these IDs so that ADC081C will again
    work + add a comment to that affect for future reference.

    Whilst here use generic firmware properties rather than the ACPI
    specific handling previously found in this driver.

    Reported-by: Kunyang Fan <Kunyang_Fan@aaeon.com.tw>
    Fixes: c458b7ca3fd0 ("iio:adc:ti-adc081c: Drop ACPI ids that seem very unlikely to be
official.")
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Cc: Andy Shevchenko <andy.shevchenko@gmail.com>
    Tested-by: Kunyang Fan <Kunyang_Fan@aaeon.com.tw> #UP-extremei11
    Link: https://lore.kernel.org/r/20211205172728.2826512-1-jic23@kernel.org
    Cc: <Stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7fb2bf469183e51027fdf0efc38b0b822af3711e
Author: Alexander Usyskin <alexander.usyskin@intel.com>
Date:   Tue Dec 28 10:20:47 2021 +0200

    mei: hbm: fix client dma reply status

    commit 6b0b80ac103b2a40c72a47c301745fd1f4ef4697 upstream.
```

```
    Don't blindly copy status value received from the firmware
    into internal client status field,
    It may be positive and ERR_PTR(ret) will translate it
    into an invalid address and the caller will crash.

    Put the error code into the client status on failure.

    Fixes: 369aea845951 ("mei: implement client dma setup.")
    Cc: <stable@vger.kernel.org> # v5.11+
    Reported-by: Emmanuel Grumbach <emmanuel.grumbach@intel.com>
    Tested-by: : Emmanuel Grumbach <emmanuel.grumbach@intel.com>
    Acked-by: Tomas Winkler <tomas.winkler@intel.com>
    Signed-off-by: Alexander Usyskin <alexander.usyskin@intel.com>
    Signed-off-by: Tomas Winkler <tomas.winkler@intel.com>
    Link: https://lore.kernel.org/r/20211228082047.378115-1-tomas.winkler@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3131efcc74cf8d98ab65e5279325e7688dad19a9
Author: Johan Hovold <johan@kernel.org>
Date:   Wed Dec 22 11:48:43 2021 +0100

    can: softing_cs: softingcs_probe(): fix memleak on registration failure

    commit ced4913efb0acc844ed65cc01d091a85d83a2082 upstream.

    In case device registration fails during probe, the driver state and
    the embedded platform device structure needs to be freed using
    platform_device_put() to properly free all resources (e.g. the device
    name).

    Fixes: 0a0b7a5f7a04 ("can: add driver for Softing card")
    Link: https://lore.kernel.org/all/20211222104843.6105-1-johan@kernel.org
    Cc: stable@vger.kernel.org # 2.6.38
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Reviewed-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5c2afdc84774acde75a9becdb8e5685d176639b3
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Wed Dec 1 13:41:25 2021 +0100

    media: cec-pin: fix interrupt en/disable handling

    commit 713bdfa10b5957053811470d298def9537d9ff13 upstream.

    The en/disable_irq() functions keep track of the 'depth': i.e. if
    interrupts are disabled twice, then it needs to enable_irq() calls to
    enable them again. The cec-pin framework didn't take this into accound
    and could disable irqs multiple times, and it expected that a single
    enable_irq() would enable them again.

    Move all calls to en/disable_irq() to the kthread where it is easy
    to keep track of the current irq state and ensure that multiple
    en/disable_irq calls never happen.

    If interrupts where disabled twice, then they would never turn on
    again, leaving the CEC adapter in a dead state.

    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Fixes: 865463fc03ed (media: cec-pin: add error injection support)
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit cb9a3f798e8c0dfda97e7a99ad220a3cef308405
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:41 2021 +0100

    media: stk1160: fix control-message timeouts

    commit 6aa6e70cdb5b863a57bad61310bf89b6617a5d2d upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: 9cb2173e6ea8 ("[media] media: Add stk1160 new driver (easycap replacement)")
    Cc: stable@vger.kernel.org      # 3.7
    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 5a7838b1f1030be46f5a38a90c1d33df6f1a4f0e
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:39 2021 +0100

    media: pvrusb2: fix control-message timeouts

    commit b82bf9b9dc305d7d3d93eab106d70dbf2171b43e upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: d855497edbfb ("V4L/DVB (4228a): pvrusb2 to kernel 2.6.18")
    Cc: stable@vger.kernel.org      # 2.6.18
    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 2e7e1c095d0024a53d3d1a21b0f4508c8e8c14c2
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:35 2021 +0100

    media: redrat3: fix control-message timeouts

    commit 2adc965c8bfa224e11ecccf9c92fd458c4236428 upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: 2154be651b90 ("[media] redrat3: new rc-core IR transceiver device driver")
    Cc: stable@vger.kernel.org      # 3.0
    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 04b45ea279c0ad404db4f21117497e8454df081b
Author: Michael Kuron <michael.kuron@gmail.com>
Date:   Sun Sep 26 21:51:26 2021 +0100

    media: dib0700: fix undefined behavior in tuner shutdown

    commit f7b77ebe6d2f49c7747b2d619586d1aa33f9ea91 upstream.

    This fixes a problem where closing the tuner would leave it in a state
    where it would not tune to any channel when reopened. This problem was
    discovered as part of https://github.com/hselasky/webcamd/issues/16.
```

```
    Since adap->id is 0 or 1, this bit-shift overflows, which is undefined
    behavior. The driver still worked in practice as the overflow would in
    most environments result in 0, which rendered the line a no-op. When
    running the driver as part of webcamd however, the overflow could lead
    to 0xff due to optimizations by the compiler, which would, in the end,
    improperly shut down the tuner.

    The bug is a regression introduced in the commit referenced below. The
    present patch causes identical behavior to before that commit for
    adap->id equal to 0 or 1. The driver does not contain support for
    dib0700 devices with more adapters, assuming such even exist.

    Tests have been performed with the Xbox One Digital TV Tuner on amd64.
    Not all dib0700 devices are expected to be affected by the regression;
    this code path is only taken by those with incorrect endpoint numbers.

    Link: https://lore.kernel.org/linux-
media/1d2fc36d94ced6f67c7cc21dcc469d5e5bdd8201.1632689033.git.mchehab+huawei@kernel.org

    Cc: stable@vger.kernel.org
    Fixes: 7757ddda6f4f ("[media] DiB0700: add function to change I2C-speed")
    Signed-off-by: Michael Kuron <michael.kuron@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 545b92aa11ac2775c2aaf9ebda148f9fec24a4a9
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:40 2021 +0100

    media: s2255: fix control-message timeouts

    commit f71d272ad4e354097020a4e6b1dc6e4b59feb50f upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Use the common control-message timeout define for the five-second
    timeouts.

    Fixes: 38f993ad8b1f ("V4L/DVB (8125): This driver adds support for the Sensoray 2255
devices.")
    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 3fca7c173e22c87bf396e6940453207a39abc95e
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:37 2021 +0100

    media: cpia2: fix control-message timeouts

    commit 10729be03327f53258cb196362015ad5c6eabe02 upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: ab33d5071de7 ("V4L/DVB (3376): Add cpia2 camera support")
    Cc: stable@vger.kernel.org      # 2.6.17
    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 6c9ffab0a81157d9f3a17c71e88d21c09732f536
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:38 2021 +0100

    media: em28xx: fix control-message timeouts

    commit d9b7e8df3aa9b8c10708aab60e72e79ac08237e4 upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: a6c2ba283565 ("[PATCH] v4l: 716: support for em28xx board family")
    Cc: stable@vger.kernel.org      # 2.6.16
    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 0bc6d86f55dcdcd755307b5905e73557df0d737e
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:34 2021 +0100

    media: mceusb: fix control-message timeouts

    commit 16394e998cbb050730536bdf7e89f5a70efbd974 upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Fixes: 66e89522aff7 ("V4L/DVB: IR: add mceusb IR receiver driver")
    Cc: stable@vger.kernel.org      # 2.6.36
    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 0bbafe22497d5cdec4bccfb0ff4c1d682dcaf43c
Author: Johan Hovold <johan@kernel.org>
Date:   Mon Oct 25 13:16:36 2021 +0100

    media: flexcop-usb: fix control-message timeouts

    commit cd1798a387825cc4a51282f5a611ad05bb1ad75f upstream.

    USB control-message timeouts are specified in milliseconds and should
    specifically not vary with CONFIG_HZ.

    Note that the driver was multiplying some of the timeout values with HZ
    twice resulting in 3000-second timeouts with HZ=1000.

    Also note that two of the timeout defines are currently unused.

    Fixes: 2154be651b90 ("[media] redrat3: new rc-core IR transceiver device driver")
    Cc: stable@vger.kernel.org      # 3.0
    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 3a1e7eb8aed9421a47ddd6a8ac54ae0b3c7ed7c4
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Wed Nov 3 12:28:31 2021 +0000

    media: v4l2-ioctl.c: readbuffers depends on V4L2_CAP_READWRITE
```

```
    commit cd9d9377ed235b294a492a094e1666178a5e78fd upstream.

    If V4L2_CAP_READWRITE is not set, then readbuffers must be set to 0,
    otherwise v4l2-compliance will complain.

    A note on the Fixes tag below: this patch does not really fix that commit,
    but it can be applied from that commit onwards. For older code there is no
    guarantee that device_caps is set, so even though this patch would apply,
    it will not work reliably.

    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Fixes: 049e684f2de9 (media: v4l2-dev: fix WARN_ON(!vdev->device_caps))
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0656b111aa6e0d9603588242c12ea3768c6c05af
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Wed Dec 15 09:38:48 2021 +0100

    media: ov8865: Disable only enabled regulators on error path

    commit cbe0b3af73bf72fad197756f026084404e2bcdc7 upstream.

    If powering on the sensor failed, the entire power-off sequence was run
    independently of how far the power-on sequence proceeded before the error.
    This lead to disabling regulators and/or clock that was not enabled.

    Fix this by disabling only clocks and regulators that were enabled
    previously.

    Fixes: 11c0d8fdccc5 ("media: i2c: Add support for the OV8865 image sensor")
    Cc: stable@vger.kernel.org
    Signed-off-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 ff56d3f33bf9b67ac620edd53bde54fc982f1a5c
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Wed Dec 1 13:41:26 2021 +0100

    media: cec: fix a deadlock situation

    commit a9e6107616bb8108aa4fc22584a05e69761a91f7 upstream.

    The cec_devnode struct has a lock meant to serialize access
    to the fields of this struct. This lock is taken during
    device node (un)registration and when opening or releasing a
    filehandle to the device node. When the last open filehandle
    is closed the cec adapter might be disabled by calling the
    adap_enable driver callback with the devnode.lock held.

    However, if during that callback a message or event arrives
    then the driver will call one of the cec_queue_event()
    variants in cec-adap.c, and those will take the same devnode.lock
    to walk the open filehandle list.

    This obviously causes a deadlock.

    This is quite easy to reproduce with the cec-gpio driver since that
    uses the cec-pin framework which generated lots of events and uses
    a kernel thread for the processing, so when adap_enable is called
    the thread is still running and can generate events.
```

```
    But I suspect that it might also happen with other drivers if an
    interrupt arrives signaling e.g. a received message before adap_enable
    had a chance to disable the interrupts.

    This patch adds a new mutex to serialize access to the fhs list.
    When adap_enable() is called the devnode.lock mutex is held, but
    not devnode.lock_fhs. The event functions in cec-adap.c will now
    use devnode.lock_fhs instead of devnode.lock, ensuring that it is
    safe to call those functions from the adap_enable callback.

    This specific issue only happens if the last open filehandle is closed
    and the physical address is invalid. This is not something that
    happens during normal operation, but it does happen when monitoring
    CEC traffic (e.g. cec-ctl --monitor) with an unconfigured CEC adapter.

    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Cc: <stable@vger.kernel.org>  # for v5.13 and up
    Signed-off-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 20f8f07ca36f316358cb13b51d21cf700543e558
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Thu Jan 6 10:30:31 2022 +0900

    ksmbd: add reserved room in ipc request/response

    commit 41dbda16a0902798e732abc6599de256b9dc3b27 upstream.

    Whenever new parameter is added to smb configuration, It is possible
    to break the execution of the IPC daemon by mismatch size of
    request/response. This patch tries to reserve space in ipc request/response
    in advance to prevent that.

    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 176b11f9497263e90df181efd6b04cc25b4cb4ab
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Fri Dec 31 09:26:25 2021 +0900

    ksmbd: limits exceeding the maximum allowable outstanding requests

    commit b589f5db6d4af8f14d70e31e1276b4c017668a26 upstream.

    If the client ignores the CreditResponse received from the server and
    continues to send the request, ksmbd limits the requests if it exceeds
    smb2 max credits.

    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4f6a52388f7ce1f154cf2d1a888114fb0bc6e102
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Wed Dec 29 23:10:03 2021 +0900

    ksmbd: move credit charge deduction under processing request

    commit 914d7e5709ac59ded70bea7956d408fe2acd7c3c upstream.

    Moves the credit charge deduction from total_credits under the processing
    a request. When repeating smb2 lock request and other command request,
    there will be a problem that ->total_credits does not decrease.
```

```
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 82be7457210c6ec7727634ac3476ad013c0a6879
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Wed Dec 29 23:08:46 2021 +0900

    ksmbd: add support for smb2 max credit parameter

    commit 004443b3f6d722b455cf963ed7c3edd7f4772405 upstream.

    Add smb2 max credits parameter to adjust maximum credits value to limit
    number of outstanding requests.

    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit be3641aea3e04f15b9ae04a2d0a2743751f38932
Author: Namjae Jeon <linkinjeon@kernel.org>
Date:   Mon Jan 17 22:16:01 2022 +0900

    ksmbd: fix guest connection failure with nautilus

    commit ac090d9c90b087d6fb714e54b2a6dd1e6c373ed6 upstream.

    MS-SMB2 describe session sign like the following.
    Session.SigningRequired MUST be set to TRUE under the following conditions:
     - If the SMB2_NEGOTIATE_SIGNING_REQUIRED bit is set in the SecurityMode
       field of the client request.
     - If the SMB2_SESSION_FLAG_IS_GUEST bit is not set in the SessionFlags
       field and Session.IsAnonymous is FALSE and either Connection.ShouldSign
       or global RequireMessageSigning is TRUE.

    When trying guest account connection using nautilus, The login failure
    happened on session setup. ksmbd does not allow this connection
    when the user is a guest and the connection sign is set. Just do not set
    session sign instead of error response as described in the specification.
    And this change improves the guest connection in Nautilus.

    Fixes: e2f34481b24d ("cifsd: add server-side procedures for SMB3")
    Cc: stable@vger.kernel.org # v5.15+
    Signed-off-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 374adb6f4746b2c06029c071086b7c1f74d980db
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Sat Jan 15 14:49:00 2022 +0300

    ksmbd: uninitialized variable in create_socket()

    commit b207602fb04537cb21ac38fabd7577eca2fa05ae upstream.

    The "ksmbd_socket" variable is not initialized on this error path.

    Cc: stable@vger.kernel.org
    Fixes: 0626e6641f6b ("cifsd: add server handler for central processing and tranport
layers")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Acked-by: Namjae Jeon <linkinjeon@kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 2bb173da2c7b8ae7f4d1ede4fd2aad989087568d
Author: Mohammad Athari Bin Ismail <mohammad.athari.ismail@intel.com>
Date:   Sat Jan 15 17:25:15 2022 +0800

    net: phy: marvell: add Marvell specific PHY loopback

    commit 020a45aff1190c32b1087cd75b57fbf6bff46ea6 upstream.

    Existing genphy_loopback() is not applicable for Marvell PHY. Besides
    configuring bit-6 and bit-13 in Page 0 Register 0 (Copper Control
    Register), it is also required to configure same bits  in Page 2
    Register 21 (MAC Specific Control Register 2) according to speed of
    the loopback is operating.

    Tested working on Marvell88E1510 PHY for all speeds (1000/100/10Mbps).

    FIXME: Based on trial and error test, it seem 1G need to have delay between
    soft reset and loopback enablement.

    Fixes: 014068dcb5b1 ("net: phy: genphy_loopback: add link speed configuration")
    Cc: <stable@vger.kernel.org> # 5.15.x
    Signed-off-by: Mohammad Athari Bin Ismail <mohammad.athari.ismail@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 01ddcec448df272b91c6250eae7c9ebe60297521
Author: Mateusz Jończyk <mat.jonczyk@o2.pl>
Date:   Fri Dec 10 21:01:23 2021 +0100

    rtc: cmos: take rtc_lock while reading from CMOS

    commit 454f47ff464325223129b9b5b8d0b61946ec704d upstream.

    Reading from the CMOS involves writing to the index register and then
    reading from the data register. Therefore access to the CMOS has to be
    serialized with rtc_lock. This invocation of CMOS_READ was not
    serialized, which could cause trouble when other code is accessing CMOS
    at the same time.

    Use spin_lock_irq() like the rest of the function.

    Nothing in kernel modifies the RTC_DM_BINARY bit, so there could be a
    separate pair of spin_lock_irq() / spin_unlock_irq() before doing the
    math.

    Signed-off-by: Mateusz Jończyk <mat.jonczyk@o2.pl>
    Reviewed-by: Nobuhiro Iwamatsu <iwamatsu@nigauri.org>
    Cc: Alessandro Zummo <a.zummo@towertech.it>
    Cc: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/20211210200131.153887-2-mat.jonczyk@o2.pl
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 09b508983f4ac4fc1afbca40bcf119a07b3c2387
Author: Willy Tarreau <w@1wt.eu>
Date:   Sun Oct 24 19:28:16 2021 +0200

    tools/nolibc: fix incorrect truncation of exit code

    commit de0244ae40ae91145faaf164a4252347607c3711 upstream.

    Ammar Faizi reported that our exit code handling is wrong. We truncate
    it to the lowest 8 bits but the syscall itself is expected to take a
    regular 32-bit signed integer, not an unsigned char. It's the kernel
```

```
    that later truncates it to the lowest 8 bits. The difference is visible
    in strace, where the program below used to show exit(255) instead of
    exit(-1):

      int main(void)
      {
            return -1;
      }

    This patch applies the fix to all archs. x86_64, i386, arm64, armv7 and
    mips were all tested and confirmed to work fine now. Risc-v was not
    tested but the change is trivial and exactly the same as for other archs.

    Reported-by: Ammar Faizi <ammar.faizi@students.amikom.ac.id>
    Cc: stable@vger.kernel.org
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 27d4111c3cf2ea093fcc466b27547b6677204be5
Author: Willy Tarreau <w@1wt.eu>
Date:   Sun Oct 24 19:28:15 2021 +0200

    tools/nolibc: i386: fix initial stack alignment

    commit ebbe0d8a449d183fa43b42d84fcb248e25303985 upstream.

    After re-checking in the spec and comparing stack offsets with glibc,
    The last pushed argument must be 16-byte aligned (i.e. aligned before the
    call) so that in the callee esp+4 is multiple of 16, so the principle is
    the 32-bit equivalent to what Ammar fixed for x86_64. It's possible that
    32-bit code using SSE2 or MMX could have been affected. In addition the
    frame pointer ought to be zero at the deepest level.

    Link: https://gitlab.com/x86-psABIs/i386-ABI/-/wikis/Intel386-psABI
    Cc: Ammar Faizi <ammar.faizi@students.amikom.ac.id>
    Cc: stable@vger.kernel.org
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b16407b920db65b90c1d829f02a5998b027e5ba5
Author: Jakub Kicinski <kuba@kernel.org>
Date:   Tue Dec 21 07:06:11 2021 -0800

    crypto: x86/aesni - don't require alignment of data

    commit d480a26bdf872529919e7c30e17f79d0d7b8c4da upstream.

    x86 AES-NI routines can deal with unaligned data. Crypto context
    (key, iv etc.) have to be aligned but we take care of that separately
    by copying it onto the stack. We were feeding unaligned data into
    crypto routines up until commit 83c83e658863 ("crypto: aesni -
    refactor scatterlist processing") switched to use the full
    skcipher API which uses cra_alignmask to decide data alignment.

    This fixes 21% performance regression in kTLS.

    Tested by booting with CONFIG_CRYPTO_MANAGER_EXTRA_TESTS=y
    (and running thru various kTLS packets).

    CC: stable@vger.kernel.org # 5.15+
    Fixes: 83c83e658863 ("crypto: aesni - refactor scatterlist processing")
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Acked-by: Ard Biesheuvel <ardb@kernel.org>
```

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

commit 38d25ad8a1ceacce30928f85fbcc2094f50bdc0a
Author: Ammar Faizi <ammar.faizi@students.amikom.ac.id>
Date:   Sun Oct 24 19:28:14 2021 +0200

    tools/nolibc: x86-64: Fix startup code bug

    commit 937ed91c712273131de6d2a02caafd3ee84e0c72 upstream.

    Before this patch, the `_start` function looks like this:
    ```
    0000000000001170 <_start>:
        1170:   pop   %rdi
        1171:   mov   %rsp,%rsi
        1174:   lea   0x8(%rsi,%rdi,8),%rdx
        1179:   and $0xfffffffffffffff0,%rsp
        117d: sub $0x8,%rsp
        1181:   call   1000 <main>
        1186:   movzbq %al,%rdi
        118a: mov $0x3c,%rax
        1191:   syscall
        1193:   hlt
        1194:   data16 cs nopw 0x0(%rax,%rax,1)
        119f: nop
    ```
    Note the "and" to %rsp with $-16, it makes the %rsp be 16-byte aligned,
    but then there is a "sub" with $0x8 which makes the %rsp no longer
    16-byte aligned, then it calls main. That's the bug!

    What actually the x86-64 System V ABI mandates is that right before the
    "call", the %rsp must be 16-byte aligned, not after the "call". So the
    "sub" with $0x8 here breaks the alignment. Remove it.

    An example where this rule matters is when the callee needs to align
    its stack at 16-byte for aligned move instruction, like `movdqa` and
    `movaps`. If the callee can't align its stack properly, it will result
    in segmentation fault.

    x86-64 System V ABI also mandates the deepest stack frame should be
    zero. Just to be safe, let's zero the %rbp on startup as the content
    of %rbp may be unspecified when the program starts. Now it looks like
    this:
    ```
    0000000000001170 <_start>:
        1170:   pop   %rdi
        1171:   mov   %rsp,%rsi
        1174:   lea   0x8(%rsi,%rdi,8),%rdx
        1179:   xor   %ebp,%ebp   # zero the %rbp
        117b: and $0xfffffffffffffff0,%rsp # align the %rsp
        117f: call   1000 <main>
        1184:   movzbq %al,%rdi
        1188:   mov $0x3c,%rax
        118f: syscall
        1191:   hlt
        1192:   data16 cs nopw 0x0(%rax,%rax,1)
        119d: nopl   (%rax)
    ```


    Cc: Bedirhan KURT <windowz414@gnuweeb.org>
    Cc: Louvian Lyndal <louvianlyndal@gmail.com>
    Reported-by: Peter Cordes <peter@cordes.ca>
    Signed-off-by: Ammar Faizi <ammar.faizi@students.amikom.ac.id>
    [wt: I did this on purpose due to a misunderstanding of the spec, other
```

```
         archs will thus have to be rechecked, particularly i386]
    Cc: stable@vger.kernel.org
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4366ccc84f1ab39f95a61a07b9679f6206ec44e0
Author: Lucas De Marchi <lucas.demarchi@intel.com>
Date:   Thu Jan 13 16:28:39 2022 -0800

    x86/gpu: Reserve stolen memory for first integrated Intel GPU

    commit 9c494ca4d3a535f9ca11ad6af1813983c1c6cbdd upstream.

    "Stolen memory" is memory set aside for use by an Intel integrated GPU.
    The intel_graphics_quirks() early quirk reserves this memory when it is
    called for a GPU that appears in the intel_early_ids[] table of integrated
    GPUs.

    Previously intel_graphics_quirks() was marked as QFLAG_APPLY_ONCE, so it
    was called only for the first Intel GPU found.  If a discrete GPU happened
    to be enumerated first, intel_graphics_quirks() was called for it but not
    for any integrated GPU found later.  Therefore, stolen memory for such an
    integrated GPU was never reserved.

    For example, this problem occurs in this Alderlake-P (integrated) + DG2
    (discrete) topology where the DG2 is found first, but stolen memory is
    associated with the integrated GPU:

      - 00:01.0 Bridge
        `- 03:00.0 DG2 discrete GPU
      - 00:02.0 Integrated GPU (with stolen memory)

    Remove the QFLAG_APPLY_ONCE flag and call intel_graphics_quirks() for every
    Intel GPU.  Reserve stolen memory for the first GPU that appears in
    intel_early_ids[].

    [bhelgaas: commit log, add code comment, squash in
    https://lore.kernel.org/r/20220118190558.2ququ4vdfjuahicm@ldmartin-desk2]
    Link: https://lore.kernel.org/r/20220114002843.2083382-1-lucas.demarchi@intel.com
    Signed-off-by: Lucas De Marchi <lucas.demarchi@intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cdee401899bac2bf847fe12d2d6b9dcd49b64ca5
Author: Jisheng Zhang <jszhang@kernel.org>
Date:   Thu Dec 2 23:36:41 2021 +0800

    riscv: mm: fix wrong phys_ram_base value for RV64

    commit b0fd4b1bf995172b9efcee23600d4f69571c321c upstream.

    Currently, if 64BIT and !XIP_KERNEL, the phys_ram_base is always 0,
    no matter the real start of dram reported by memblock is.

    Fixes: 6d7f91d914bc ("riscv: Get rid of CONFIG_PHYS_RAM_BASE in kernel physical address
conversion")
    Signed-off-by: Jisheng Zhang <jszhang@kernel.org>
    Reviewed-by: Alexandre Ghiti <alex@ghiti.fr>
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 128aeb0b5161cad2447ff101f8c6127b52dafa87
```

```
Author: Nick Kossifidis <mick@ics.forth.gr>
Date:   Fri Nov 26 20:04:10 2021 +0200

    riscv: use hart id instead of cpu id on machine_kexec

    commit 0e105f1d0037d677dff3c697d22f9551e6c39af8 upstream.

    raw_smp_processor_id() doesn't return the hart id as stated in
    arch/riscv/include/asm/smp.h, use smp_processor_id() instead
    to get the cpu id, and cpuid_to_hartid_map() to pass the hart id
    to the next kernel. This fixes kexec on HiFive Unleashed/Unmatched
    where cpu ids and hart ids don't match (on qemu-virt they match).

    Fixes: fba8a8674f68 ("RISC-V: Add kexec support")
    Signed-off-by: Nick Kossifidis <mick@ics.forth.gr>
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 078e59be6f04bf43eb7e842d7bdb9dea73a312a7
Author: Nick Kossifidis <mick@ics.forth.gr>
Date:   Fri Nov 26 20:04:09 2021 +0200

    riscv: Don't use va_pa_offset on kdump

    commit a11c07f032a0e9a562a32ece73af96b0e754c4b3 upstream.

    On kdump instead of using an intermediate step to relocate the kernel,
    that lives in a "control buffer" outside the current kernel's mapping,
    we jump to the crash kernel directly by calling riscv_kexec_norelocate().
    The current implementation uses va_pa_offset while switching to physical
    addressing, however since we moved the kernel outside the linear mapping
    this won't work anymore since riscv_kexec_norelocate() is part of the
    kernel mapping and we should use kernel_map.va_kernel_pa_offset, and also
    take XIP kernel into account.

    We don't really need to use va_pa_offset on riscv_kexec_norelocate, we
    can just set STVEC to the physical address of the new kernel instead and
    let the hart jump to the new kernel on the next instruction after setting
    SATP to zero. This fixes kdump and is also simpler/cleaner.

    I tested this on the latest qemu and HiFive Unmatched and works as
    expected.

    Fixes: 2bfc6cd81bd1 ("riscv: Move kernel mapping outside of linear mapping")
    Signed-off-by: Nick Kossifidis <mick@ics.forth.gr>
    Reviewed-by: Alexandre Ghiti <alex@ghiti.fr>
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7d077c2b2bba10306cc5929bc81e0721ec8b5374
Author: Nick Kossifidis <mick@ics.forth.gr>
Date:   Fri Nov 26 20:04:11 2021 +0200

    riscv: try to allocate crashkern region from 32bit addressible memory

    commit decf89f86ecd3c3c3de81c562010d5797bea3de1 upstream.

    When allocating crash kernel region without explicitly specifying its
    base address/size, memblock_phys_alloc_range will attempt to allocate
    memory top to bottom (memblock.bottom_up is false), so the crash
    kernel region will end up in highmem on 64bit systems. This way
    swiotlb can't work on the crash kernel, since there won't be any
    32bit addressible memory available for the bounce buffers.
```

```
    Try to allocate 32bit addressible memory if available, for the
    crash kernel by restricting the top search address to be less
    than SZ_4G. If that fails fallback to the previous behavior.

    I tested this on HiFive Unmatched where the pci-e controller needs
    swiotlb to work, with this patch it's possible to access the pci-e
    controller on crash kernel and mount the rootfs from the nvme.

    Signed-off-by: Nick Kossifidis <mick@ics.forth.gr>
    Fixes: e53d28180d4d ("RISC-V: Add kdump support")
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9d213b77f85316b69ab472ddd195f68003ebc48b
Author: Sean Christopherson <seanjc@google.com>
Date:   Mon Nov 29 21:43:42 2021 +0000

    RISC-V: Use common riscv_cpuid_to_hartid_mask() for both SMP=y and SMP=n

    commit 869c70609248102f3a2e95a39b6233ff6ea2c932 upstream.

    Use what is currently the SMP=y version of riscv_cpuid_to_hartid_mask()
    for both SMP=y and SMP=n to fix a build failure with KVM=m and SMP=n due
    to boot_cpu_hartid not being exported.  This also fixes a second bug
    where the SMP=n version assumes the sole CPU in the system is in the
    incoming mask, which may not hold true in kvm_riscv_vcpu_sbi_ecall() if
    the KVM guest VM has multiple vCPUs (on a SMP=n system).

    Fixes: 1ef46c231df4 ("RISC-V: Implement new SBI v0.2 extensions")
    Reported-by: Adam Borowski <kilobyte@angband.pl>
    Reviewed-by: Anup Patel <anup.patel@wdc.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5b0cc071e980467e28f94d642de3e8656fece31c
Author: Alexandre Ghiti <alexandre.ghiti@canonical.com>
Date:   Mon Jan 17 10:57:16 2022 +0100

    riscv: Get rid of MAXPHYSMEM configs

    commit db1503d355a79d1d4255a9996f20e72848b74a56 upstream.

    CONFIG_MAXPHYSMEM_* are actually never used, even the nommu defconfigs
    selecting the MAXPHYSMEM_2GB had no effects on PAGE_OFFSET since it was
    preempted by !MMU case right before.

    In addition, the move of the kernel mapping at the end of the address
    space broke the use of MAXPHYSMEM_2G with MMU since it defines PAGE_OFFSET
    at the same address as the kernel mapping.

    Reported-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Fixes: 2bfc6cd81bd1 ("riscv: Move kernel mapping outside of linear mapping")
    Signed-off-by: Alexandre Ghiti <alexandre.ghiti@canonical.com>
    Tested-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Tested-by: Conor Dooley <Conor.Dooley@microchip.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 627d769bb1fc787bd01a849e346f173a5fea34d7
Author: Paul Cercueil <paul@crapouillou.net>
```

```
Date:   Sat Oct 16 14:22:28 2021 +0100

    mtd: rawnand: ingenic: JZ4740 needs 'oob_first' read page function

    commit 0171480007d64f663aae9226303f1b1e4621229e upstream.

    The ECC engine on the JZ4740 SoC requires the ECC data to be read before
    the page; using the default page reading function does not work. Indeed,
    the old JZ4740 NAND driver (removed in 5.4) did use the 'OOB first' flag
    that existed back then.

    Use the newly created nand_read_page_hwecc_oob_first() to address this
    issue.

    This issue was not found when the new ingenic-nand driver was developed,
    most likely because the Device Tree used had the nand-ecc-mode set to
    "hw_oob_first", which seems to not be supported anymore.

    Cc: <stable@vger.kernel.org> # v5.2
    Fixes: a0ac778eb82c ("mtd: rawnand: ingenic: Add support for the JZ4740")
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211016132228.40254-5-paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 906261d4e19700cbda23cb1c9d38098b7c75178f
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sat Oct 16 14:22:27 2021 +0100

    mtd: rawnand: Export nand_read_page_hwecc_oob_first()

    commit d8466f73010faf71effb21228ae1cbf577dab130 upstream.

    Move the function nand_read_page_hwecc_oob_first() (previously
    nand_davinci_read_page_hwecc_oob_first()) to nand_base.c, and export it
    as a GPL symbol, so that it can be used by more modules.

    Cc: <stable@vger.kernel.org> # v5.2
    Fixes: a0ac778eb82c ("mtd: rawnand: ingenic: Add support for the JZ4740")
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211016132228.40254-4-paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2c7bc80d7c78438940dde0ad98bf0e8bb2e76053
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sat Oct 16 14:22:26 2021 +0100

    mtd: rawnand: davinci: Rewrite function description

    commit 0697f8441faad552fbeb02d74454b5e7bcc956a2 upstream.

    The original comment that describes the function
    nand_davinci_read_page_hwecc_oob_first() is very obscure and it is hard
    to understand what it is for.

    Cc: <stable@vger.kernel.org> # v5.2
    Fixes: a0ac778eb82c ("mtd: rawnand: ingenic: Add support for the JZ4740")
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211016132228.40254-3-paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 70fc7db53e537387020f852ad0a7adb0a93c1241
Author: Paul Cercueil <paul@crapouillou.net>
```

```
Date:   Sat Oct 16 14:22:25 2021 +0100

    mtd: rawnand: davinci: Avoid duplicated page read

    commit 9c9d709965385de5a99f84b14bd5860e1541729e upstream.

    The function nand_davinci_read_page_hwecc_oob_first() first reads the
    OOB data, extracts the ECC information, programs the ECC hardware before
    reading the actual data in a loop.

    Right after the OOB data was read, it called nand_read_page_op() to
    reset the read cursor to the beginning of the page. This caused the
    first page to be read twice: in that call, and later in the loop.

    Address that issue by changing the call to nand_read_page_op() to
    nand_change_read_column_op(), which will only reset the read cursor.

    Cc: <stable@vger.kernel.org> # v5.2
    Fixes: a0ac778eb82c ("mtd: rawnand: ingenic: Add support for the JZ4740")
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211016132228.40254-2-paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f024ac85d9cc1649c2d0bf79b50b3489ed2e4620
Author: Paul Cercueil <paul@crapouillou.net>
Date:   Sat Oct 16 14:22:24 2021 +0100

    mtd: rawnand: davinci: Don't calculate ECC when reading page

    commit 71e89591502d737c10db2bd4d8fcfaa352552afb upstream.

    The function nand_davinci_read_page_hwecc_oob_first() does read the ECC
    data from the OOB area. Therefore it does not need to calculate the ECC
    as it is already available.

    Cc: <stable@vger.kernel.org> # v5.2
    Fixes: a0ac778eb82c ("mtd: rawnand: ingenic: Add support for the JZ4740")
    Signed-off-by: Paul Cercueil <paul@crapouillou.net>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211016132228.40254-1-paul@crapouillou.net
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cf4534a455bed07ad817ca22240ee94ec1a207c6
Author: Andreas Oetken <ennoerlangen@gmail.com>
Date:   Tue Nov 2 18:26:04 2021 +0100

    mtd: Fixed breaking list in __mtd_del_partition.

    commit 2966daf7d253d9904b337b040dd7a43472858b8a upstream.

    Not the child partition should be removed from the partition list
    but the partition itself. Otherwise the partition list gets broken
    and any subsequent remove operations leads to a kernel panic.

    Fixes: 46b5889cc2c5 ("mtd: implement proper partition handling")
    Signed-off-by: Andreas Oetken <andreas.oetken@siemens-energy.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211102172604.2921065-1-andreas.oetken@siemens-
energy.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb9366b717b56a3bb60100533595866b863f3450
Author: Stefan Riedmueller <s.riedmueller@phytec.de>
```

```
Date:   Tue Nov 2 21:20:21 2021 +0100

    mtd: rawnand: gpmi: Remove explicit default gpmi clock setting for i.MX6

    commit aa1baa0e6c1aa4872e481dce4fc7fd6f3dd8496b upstream.

    There is no need to explicitly set the default gpmi clock rate during
    boot for the i.MX 6 since this is done during nand_detect anyway.

    Signed-off-by: Stefan Riedmueller <s.riedmueller@phytec.de>
    Cc: stable@vger.kernel.org
    Acked-by: Han Xu <han.xu@nxp.com>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211102202022.15551-1-ceggers@arri.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c447696e2f825df7800b0630352bea2d45d09baa
Author: Christian Eggers <ceggers@arri.de>
Date:   Tue Nov 2 21:20:22 2021 +0100

    mtd: rawnand: gpmi: Add ERR007117 protection for nfc_apply_timings

    commit f53d4c109a666bf1a4883b45d546fba079258717 upstream.

    gpmi_io clock needs to be gated off when changing the parent/dividers of
    enfc_clk_root (i.MX6Q/i.MX6UL) respectively qspi2_clk_root (i.MX6SX).
    Otherwise this rate change can lead to an unresponsive GPMI core which
    results in DMA timeouts and failed driver probe:

    [    4.072318] gpmi-nand 112000.gpmi-nand: DMA timeout, last DMA
    ...
    [    4.370355] gpmi-nand 112000.gpmi-nand: Chip: 0, Error -110
    ...
    [    4.375988] gpmi-nand 112000.gpmi-nand: Chip: 0, Error -22
    [    4.381524] gpmi-nand 112000.gpmi-nand: Error in ECC-based read: -22
    [    4.387988] gpmi-nand 112000.gpmi-nand: Chip: 0, Error -22
    [    4.393535] gpmi-nand 112000.gpmi-nand: Chip: 0, Error -22
    ...

    Other than stated in i.MX 6 erratum ERR007117, it should be sufficient
    to gate only gpmi_io because all other bch/nand clocks are derived from
    different clock roots.

    The i.MX6 reference manuals state that changing clock muxers can cause
    glitches but are silent about changing dividers. But tests showed that
    these glitches can definitely happen on i.MX6ULL. For i.MX7D/8MM in turn,
    the manual guarantees that no glitches can happen when changing
    dividers.

    Co-developed-by: Stefan Riedmueller <s.riedmueller@phytec.de>
    Signed-off-by: Stefan Riedmueller <s.riedmueller@phytec.de>
    Signed-off-by: Christian Eggers <ceggers@arri.de>
    Cc: stable@vger.kernel.org
    Acked-by: Han Xu <han.xu@nxp.com>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20211102202022.15551-2-ceggers@arri.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c915786cb05a69b24dc85051b6bd2cd5a5f929b5
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Wed Jan 19 08:48:16 2022 +0100

    nfc: llcp: fix NULL error pointer dereference on sendmsg() after failed bind()

    commit dded08927ca3c31a5c37f8e7f95fe98770475dd4 upstream.
```

```
    Syzbot detected a NULL pointer dereference of nfc_llcp_sock->dev pointer
    (which is a 'struct nfc_dev *') with calls to llcp_sock_sendmsg() after
    a failed llcp_sock_bind(). The message being sent is a SOCK_DGRAM.

    KASAN report:

      BUG: KASAN: null-ptr-deref in nfc_alloc_send_skb+0x2d/0xc0
      Read of size 4 at addr 00000000000005c8 by task llcp_sock_nfc_a/899

      CPU: 5 PID: 899 Comm: llcp_sock_nfc_a Not tainted 5.16.0-rc6-next-20211224-00001-
gc6437fbf18b0 #125
      Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.14.0-2 04/01/2014
      Call Trace:
       <TASK>
       dump_stack_lvl+0x45/0x59
       ? nfc_alloc_send_skb+0x2d/0xc0
       __kasan_report.cold+0x117/0x11c
       ? mark_lock+0x480/0x4f0
       ? nfc_alloc_send_skb+0x2d/0xc0
       kasan_report+0x38/0x50
       nfc_alloc_send_skb+0x2d/0xc0
       nfc_llcp_send_ui_frame+0x18c/0x2a0
       ? nfc_llcp_send_i_frame+0x230/0x230
       ? __local_bh_enable_ip+0x86/0xe0
       ? llcp_sock_connect+0x470/0x470
       ? llcp_sock_connect+0x470/0x470
       sock_sendmsg+0x8e/0xa0
       ____sys_sendmsg+0x253/0x3f0
       ...

    The issue was visible only with multiple simultaneous calls to bind() and
    sendmsg(), which resulted in most of the bind() calls to fail.  The
    bind() was failing on checking if there is available WKS/SDP/SAP
    (respective bit in 'struct nfc_llcp_local' fields).  When there was no
    available WKS/SDP/SAP, the bind returned error but the sendmsg() to such
    socket was able to trigger mentioned NULL pointer dereference of
    nfc_llcp_sock->dev.

    The code looks simply racy and currently it protects several paths
    against race with checks for (!nfc_llcp_sock->local) which is NULL-ified
    in error paths of bind().  The llcp_sock_sendmsg() did not have such
    check but called function nfc_llcp_send_ui_frame() had, although not
    protected with lock_sock().

    Therefore the race could look like (same socket is used all the time):
      CPU0                                      CPU1
      ====                                      ====
      llcp_sock_bind()
      - lock_sock()
        - success
      - release_sock()
      - return 0
                                                llcp_sock_sendmsg()
                                                - lock_sock()
                                                - release_sock()
      llcp_sock_bind(), same socket
      - lock_sock()
        - error
                                                - nfc_llcp_send_ui_frame()
                                                  - if (!llcp_sock->local)
        - llcp_sock->local = NULL
        - nfc_put_device(dev)
                                                    - dereference llcp_sock->dev
      - release_sock()
```

```
        - return -ERRNO

      The nfc_llcp_send_ui_frame() checked llcp_sock->local outside of the
      lock, which is racy and ineffective check.  Instead, its caller
      llcp_sock_sendmsg(), should perform the check inside lock_sock().

      Reported-and-tested-by: syzbot+7f23bcddf626e0593a39@syzkaller.appspotmail.com
      Fixes: b874dec21d1c ("NFC: Implement LLCP connection less Tx path")
      Cc: <stable@vger.kernel.org>
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 579010e18a703deafb3be95a5292a110381f92c2
  Author: Jaegeuk Kim <jaegeuk@kernel.org>
  Date:   Thu Dec 9 10:25:43 2021 -0800

      f2fs: avoid EINVAL by SBI_NEED_FSCK when pinning a file

      commit 19bdba5265624ba6b9d9dd936a0c6ccc167cfe80 upstream.

      Android OTA failed due to SBI_NEED_FSCK flag when pinning the file. Let's avoid
      it since we can do in-place-updates.

      Cc: stable@vger.kernel.org
      Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 258b26a34778cde43f228a392e242d3d0420624a
  Author: Chao Yu <chao@kernel.org>
  Date:   Sun Dec 12 17:16:30 2021 +0800

      f2fs: fix to do sanity check on last xattr entry in __f2fs_setxattr()

      commit 645a3c40ca3d40cc32b4b5972bf2620f2eb5dba6 upstream.

      As Wenqing Liu reported in bugzilla:

      https://bugzilla.kernel.org/show_bug.cgi?id=215235

      - Overview
      page fault in f2fs_setxattr() when mount and operate on corrupted image

      - Reproduce
      tested on kernel 5.16-rc3, 5.15.X under root

      1. unzip tmp7.zip
      2. ./single.sh f2fs 7

      Sometimes need to run the script several times

      - Kernel dump
      loop0: detected capacity change from 0 to 131072
      F2FS-fs (loop0): Found nat_bits in checkpoint
      F2FS-fs (loop0): Mounted with checkpoint version = 7548c2ee
      BUG: unable to handle page fault for address: ffffe47bc7123f48
      RIP: 0010:kfree+0x66/0x320
      Call Trace:
       __f2fs_setxattr+0x2aa/0xc00 [f2fs]
       f2fs_setxattr+0xfa/0x480 [f2fs]
       __f2fs_set_acl+0x19b/0x330 [f2fs]
       __vfs_removexattr+0x52/0x70
       __vfs_removexattr_locked+0xb1/0x140
       vfs_removexattr+0x56/0x100
       removexattr+0x57/0x80
```

```
      path_removexattr+0xa3/0xc0
      __x64_sys_removexattr+0x17/0x20
      do_syscall_64+0x37/0xb0
      entry_SYSCALL_64_after_hwframe+0x44/0xae

    The root cause is in __f2fs_setxattr(), we missed to do sanity check on
    last xattr entry, result in out-of-bound memory access during updating
    inconsistent xattr data of target inode.

    After the fix, it can detect such xattr inconsistency as below:

    F2FS-fs (loop11): inode (7) has invalid last xattr entry, entry_size: 60676
    F2FS-fs (loop11): inode (8) has corrupted xattr
    F2FS-fs (loop11): inode (8) has corrupted xattr
    F2FS-fs (loop11): inode (8) has invalid last xattr entry, entry_size: 47736

    Cc: stable@vger.kernel.org
    Reported-by: Wenqing Liu <wenqingliu0120@gmail.com>
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 25c94a4fb3c5abf427fb3ca89127cb867954ac7e
Author: Chao Yu <chao@kernel.org>
Date:   Mon Dec 6 22:44:21 2021 +0800

    f2fs: fix to do sanity check in is_alive()

    commit 77900c45ee5cd5da63bd4d818a41dbdf367e81cd upstream.

    In fuzzed image, SSA table may indicate that a data block belongs to
    invalid node, which node ID is out-of-range (0, 1, 2 or max_nid), in
    order to avoid migrating inconsistent data in such corrupted image,
    let's do sanity check anyway before data block migration.

    Cc: stable@vger.kernel.org
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d667b9f61df7bdfcb59dd1406fd2392c358f0008
Author: Chao Yu <chao@kernel.org>
Date:   Mon Dec 6 22:44:19 2021 +0800

    f2fs: fix to do sanity check on inode type during garbage collection

    commit 9056d6489f5a41cfbb67f719d2c0ce61ead72d9f upstream.

    As report by Wenqing Liu in bugzilla:

    https://bugzilla.kernel.org/show_bug.cgi?id=215231

    - Overview
    kernel NULL pointer dereference triggered  in folio_mark_dirty() when mount and operate
on a crafted f2fs image

    - Reproduce
    tested on kernel 5.16-rc3, 5.15.X under root

    1. mkdir mnt
    2. mount -t f2fs tmp1.img mnt
    3. touch tmp
    4. cp tmp mnt

    F2FS-fs (loop0): sanity_check_inode: inode (ino=49) extent info [5942, 4294180864, 4] is
```

```
 incorrect, run fsck to fix
     F2FS-fs (loop0): f2fs_check_nid_range: out-of-range nid=31340049, run fsck to fix.
     BUG: kernel NULL pointer dereference, address: 0000000000000000
      folio_mark_dirty+0x33/0x50
      move_data_page+0x2dd/0x460 [f2fs]
      do_garbage_collect+0xc18/0x16a0 [f2fs]
      f2fs_gc+0x1d3/0xd90 [f2fs]
      f2fs_balance_fs+0x13a/0x570 [f2fs]
      f2fs_create+0x285/0x840 [f2fs]
      path_openat+0xe6d/0x1040
      do_filp_open+0xc5/0x140
      do_sys_openat2+0x23a/0x310
      do_sys_open+0x57/0x80

     The root cause is for special file: e.g. character, block, fifo or socket file,
     f2fs doesn't assign address space operations pointer array for mapping->a_ops field,
     so, in a fuzzed image, SSA table indicates a data block belong to special file, when
     f2fs tries to migrate that block, it causes NULL pointer access once move_data_page()
     calls a_ops->set_dirty_page().

     Cc: stable@vger.kernel.org
     Reported-by: Wenqing Liu <wenqingliu0120@gmail.com>
     Signed-off-by: Chao Yu <chao@kernel.org>
     Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 19d1540fbb5592889e182df0499b8286518553e6
 Author: Shyam Prasad N <sprasad@microsoft.com>
 Date:   Mon Jan 17 00:20:47 2022 -0600

     cifs: free ntlmsspblob allocated in negotiate

     commit e3548aaf41a200c2af359462be23bcdd76efd795 upstream.

     One of my previous fixes:
     cifs: send workstation name during ntlmssp session setup
     ...changed the prototype of build_ntlmssp_negotiate_blob
     from being allocated by the caller to being allocated within
     the function. The caller needs to free this object too.
     While SMB2 version of the caller did it, I forgot to free
     for the SMB1 version. Fixing that here.

     Fixes: 49bd49f983b5 ("cifs: send workstation name during ntlmssp session setup")
     Cc: stable@vger.kernel.org # 5.16
     Signed-off-by: Shyam Prasad N <sprasad@microsoft.com>
     Signed-off-by: Steve French <stfrench@microsoft.com>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 696e08c56d20aed84f64e0bf67ce607fc535fbeb
 Author: Takashi Iwai <tiwai@suse.de>
 Date:   Sun Jan 16 09:28:38 2022 +0100

     ALSA: core: Fix SSID quirk lookup for subvendor=0

     commit 5576c4f24c56722a2d9fb9c447d896e5b312078b upstream.

     Some weird devices set the codec SSID vendor ID 0, and
     snd_pci_quirk_lookup_id() loop aborts at the point although it should
     still try matching with the SSID device ID.  This resulted in a
     missing quirk for some old Macs.

     Fix the loop termination condition to check both subvendor and
     subdevice.

     Fixes: 73355ddd8775 ("ALSA: hda: Code refactoring snd_hda_pick_fixup()")
```

```
    Cc: <stable@vger.kernel.org>
    BugLink: https://bugzilla.kernel.org/show_bug.cgi?id=215495
    Link: https://lore.kernel.org/r/20220116082838.19382-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 683af390ae3342810d75d2799c43f0c829cac4b3
Author: Jason Gerecke <killertofu@gmail.com>
Date:   Tue Jan 18 14:38:41 2022 -0800

    HID: wacom: Avoid using stale array indicies to read contact count

    commit 20f3cf5f860f9f267a6a6e5642d3d0525edb1814 upstream.

    If we ever see a touch report with contact count data we initialize
    several variables used to read the contact count in the pre-report
    phase. These variables are never reset if we process a report which
    doesn't contain a contact count, however. This can cause the pre-
    report function to trigger a read of arbitrary memory (e.g. NULL
    if we're lucky) and potentially crash the driver.

    This commit restores resetting of the variables back to default
    "none" values that were used prior to the commit mentioned
    below.

    Link: https://github.com/linuxwacom/input-wacom/issues/276
    Fixes: 003f50ab673c (HID: wacom: Update last_slot_field during pre_report phase)
    CC: stable@vger.kernel.org
    Signed-off-by: Jason Gerecke <jason.gerecke@wacom.com>
    Reviewed-by: Ping Cheng <ping.cheng@wacom.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6414c021a0c224600600902eacbd4139761bdc2b
Author: Jason Gerecke <killertofu@gmail.com>
Date:   Tue Jan 18 14:37:56 2022 -0800

    HID: wacom: Ignore the confidence flag when a touch is removed

    commit df03e9bd6d4806619b4cdc91a3d7695818a8e2b7 upstream.

    AES hardware may internally re-classify a contact that it thought was
    intentional as a palm. Intentional contacts are reported as "down" with
    the confidence bit set. When this re-classification occurs, however, the
    state transitions to "up" with the confidence bit cleared. This kind of
    transition appears to be legal according to Microsoft docs, but we do
    not handle it correctly. Because the confidence bit is clear, we don't
    call `wacom_wac_finger_slot` and update userspace. This causes hung
    touches that confuse userspace and interfere with pen arbitration.

    This commit adds a special case to ignore the confidence flag if a contact
    is reported as removed. This ensures we do not leave a hung touch if one
    of these re-classification events occured. Ideally we'd have some way to
    also let userspace know that the touch has been re-classified as a palm
    and needs to be canceled, but that's not possible right now :)

    Link: https://github.com/linuxwacom/input-wacom/issues/288
    Fixes: 7fb0413baa7f (HID: wacom: Use "Confidence" flag to prevent reporting invalid
contacts)
    CC: stable@vger.kernel.org
    Signed-off-by: Jason Gerecke <jason.gerecke@wacom.com>
    Reviewed-by: Ping Cheng <ping.cheng@wacom.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit ea51421ab066dde9c8367e42d93dfcc6472dacaa
Author: Jason Gerecke <killertofu@gmail.com>
Date:   Tue Jan 18 14:37:55 2022 -0800

    HID: wacom: Reset expected and received contact counts at the same time

    commit 546e41ac994cc185ef3de610ca849a294b5df3ba upstream.

    These two values go hand-in-hand and must be valid for the driver to
    behave correctly. We are currently lazy about updating the values and
    rely on the "expected" code flow to take care of making sure they're
    valid at the point they're needed. The "expected" flow changed somewhat
    with commit f8b6a74719b5 ("HID: wacom: generic: Support multiple tools
    per report"), however. This led to problems with the DTH-2452 due (in
    part) to *all* contacts being fully processed -- even those past the
    expected contact count. Specifically, the received count gets reset to
    0 once all expected fingers are processed, but not the expected count.
    The rest of the contacts in the report are then *also* processed since
    now the driver thinks we've only processed 0 of N expected contacts.

    Later commits such as 7fb0413baa7f (HID: wacom: Use "Confidence" flag to
    prevent reporting invalid contacts) worked around the DTH-2452 issue by
    skipping the invalid contacts at the end of the report, but this is not
    a complete fix. The confidence flag cannot be relied on when a contact
    is removed (see the following patch), and dealing with that condition
    re-introduces the DTH-2452 issue unless we also address this contact
    count laziness. By resetting expected and received counts at the same
    time we ensure the driver understands that there are 0 more contacts
    expected in the report. Similarly, we also make sure to reset the
    received count if for some reason we're out of sync in the pre-report
    phase.

    Link: https://github.com/linuxwacom/input-wacom/issues/288
    Fixes: f8b6a74719b5 ("HID: wacom: generic: Support multiple tools per report")
    CC: stable@vger.kernel.org
    Signed-off-by: Jason Gerecke <jason.gerecke@wacom.com>
    Reviewed-by: Ping Cheng <ping.cheng@wacom.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 09874875ebaedfa0024ce0ab4ad7b3bfb6ca6ee2
Author: Jann Horn <jannh@google.com>
Date:   Fri Jan 14 14:33:30 2022 +0100

    HID: uhid: Fix worker destroying device without any protection

    commit 4ea5763fb79ed89b3bdad455ebf3f33416a81624 upstream.

    uhid has to run hid_add_device() from workqueue context while allowing
    parallel use of the userspace API (which is protected with ->devlock).
    But hid_add_device() can fail. Currently, that is handled by immediately
    destroying the associated HID device, without using ->devlock - but if
    there are concurrent requests from userspace, that's wrong and leads to
    NULL dereferences and/or memory corruption (via use-after-free).

    Fix it by leaving the HID device as-is in the worker. We can clean it up
    later, either in the UHID_DESTROY command handler or in the ->release()
    handler.

    Cc: stable@vger.kernel.org
    Fixes: 67f8ecc550b5 ("HID: uhid: fix timeout when probe races with IO")
    Signed-off-by: Jann Horn <jannh@google.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit bf1ce5847e2b885b8c37cd84c2b1589ddf33a98d
Author: Karl Kurbjun <kkurbjun@gmail.com>
Date:   Sun Jan 9 20:49:35 2022 -0700

    HID: Ignore battery for Elan touchscreen on HP Envy X360 15t-dr100

    commit f3193ea1b6779023334faa72b214ece457e02656 upstream.

    Battery status on Elan tablet driver is reported for the HP ENVY x360
    15t-dr100. There is no separate battery for the Elan controller resulting in a
    battery level report of 0% or 1% depending on whether a stylus has interacted
    with the screen. These low battery level reports causes a variety of bad
    behavior in desktop environments. This patch adds the appropriate quirk to
    indicate that the batery status is unused for this target.

    Cc: stable@vger.kernel.org
    Signed-off-by: Karl Kurbjun <kkurbjun@gmail.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4c5ccf948f867f311daec9ab20d479b4466e7cd0
Author: Marcelo Tosatti <mtosatti@redhat.com>
Date:   Tue Jan 18 04:34:43 2022 -0500

    KVM: VMX: switch blocked_vcpu_on_cpu_lock to raw spinlock

    commit 5f02ef741a785678930f3ff0a8b6b2b0ef1bb402 upstream.

    blocked_vcpu_on_cpu_lock is taken from hard interrupt context
    (pi_wakeup_handler), therefore it cannot sleep.

    Switch it to a raw spinlock.

    Fixes:

    [41297.066254] BUG: scheduling while atomic: CPU 0/KVM/635218/0x00010001
    [41297.066323] Preemption disabled at:
    [41297.066324] [<ffffffff902ee47f>] irq_enter_rcu+0xf/0x60
    [41297.066339] Call Trace:
    [41297.066342]  <IRQ>
    [41297.066346]  dump_stack_lvl+0x34/0x44
    [41297.066353]  ? irq_enter_rcu+0xf/0x60
    [41297.066356]  __schedule_bug.cold+0x7d/0x8b
    [41297.066361]  __schedule+0x439/0x5b0
    [41297.066365]  ? task_blocks_on_rt_mutex.constprop.0.isra.0+0x1b0/0x440
    [41297.066369]  schedule_rtlock+0x1e/0x40
    [41297.066371]  rtlock_slowlock_locked+0xf1/0x260
    [41297.066374]  rt_spin_lock+0x3b/0x60
    [41297.066378]  pi_wakeup_handler+0x31/0x90 [kvm_intel]
    [41297.066388]  sysvec_kvm_posted_intr_wakeup_ipi+0x9d/0xd0
    [41297.066392]  </IRQ>
    [41297.066392]  asm_sysvec_kvm_posted_intr_wakeup_ipi+0x12/0x20
    ...

    Signed-off-by: Marcelo Tosatti <mtosatti@redhat.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1259fb5a830d4341bce46fbf0d58931927f047c8
Author: David Matlack <dmatlack@google.com>
Date:   Thu Jan 13 23:30:17 2022 +0000

    KVM: x86/mmu: Fix write-protection of PTs mapped by the TDP MMU
```

commit 7c8a4742c4abe205ec9daf416c9d42fd6b406e8e upstream.

When the TDP MMU is write-protection GFNs for page table protection (as
opposed to for dirty logging, or due to the HVA not being writable), it
checks if the SPTE is already write-protected and if so skips modifying
the SPTE and the TLB flush.

This behavior is incorrect because it fails to check if the SPTE
is write-protected for page table protection, i.e. fails to check
that MMU-writable is '0'. If the SPTE was write-protected for dirty
logging but not page table protection, the SPTE could locklessly be made
writable, and vCPUs could still be running with writable mappings cached
in their TLB.

Fix this by only skipping setting the SPTE if the SPTE is already
write-protected *and* MMU-writable is already clear. Technically,
checking only MMU-writable would suffice; a SPTE cannot be writable
without MMU-writable being set. But check both to be paranoid and
because it arguably yields more readable code.

Fixes: 46044f72c382 ("kvm: x86/mmu: Support write protection for nesting in tdp MMU")
Cc: stable@vger.kernel.org
Signed-off-by: David Matlack <dmatlack@google.com>
Message-Id: <20220113233020.3986005-2-dmatlack@google.com>
Reviewed-by: Sean Christopherson <seanjc@google.com>
Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>