```
commit d7f82a4f24cb99ab82e6782567c87bd14b77d166
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Wed Jun 16 12:05:17 2021 +0200

    Linux 5.12.11

    Link: https://lore.kernel.org/r/20210614102658.137943264@linuxfoundation.org
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Jason Self <jason@bluehome.net>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9be4a3eba8c9156a4447654d6a5e69bd883e573a
Author: Linus Torvalds <torvalds@linux-foundation.org>
Date:   Tue Jun 15 09:26:19 2021 -0700

    proc: only require mm_struct for writing

    commit 94f0b2d4a1d0c52035aef425da5e022bd2cb1c71 upstream.

    Commit 591a22c14d3f ("proc: Track /proc/$pid/attr/ opener mm_struct") we
    started using __mem_open() to track the mm_struct at open-time, so that
    we could then check it for writes.

    But that also ended up making the permission checks at open time much
    stricter - and not just for writes, but for reads too.  And that in turn
    caused a regression for at least Fedora 29, where NIC interfaces fail to
    start when using NetworkManager.

    Since only the write side wanted the mm_struct test, ignore any failures
    by __mem_open() at open time, leaving reads unaffected.  The write()
    time verification of the mm_struct pointer will then catch the failure
    case because a NULL pointer will not match a valid 'current->mm'.

    Link: https://lore.kernel.org/netdev/YMjTlp2FSJYvoyFa@unreal/
    Fixes: 591a22c14d3f ("proc: Track /proc/$pid/attr/ opener mm_struct")
    Reported-and-tested-by: Leon Romanovsky <leon@kernel.org>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Christian Brauner <christian.brauner@ubuntu.com>
    Cc: Andrea Righi <andrea.righi@canonical.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 34831232890dfed49dec506638addb7938e47b2b
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Jun 2 21:30:29 2021 +0800

    scsi: core: Only put parent device if host state differs from SHOST_CREATED

    commit 1e0d4e6225996f05271de1ebcb1a7c9381af0b27 upstream.

    get_device(shost->shost_gendev.parent) is called after host state has
    switched to SHOST_RUNNING. scsi_host_dev_release() shouldn't release the
    parent device if host state is still SHOST_CREATED.

    Link: https://lore.kernel.org/r/20210602133029.2864069-5-ming.lei@redhat.com
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: John Garry <john.garry@huawei.com>
    Cc: Hannes Reinecke <hare@suse.de>
```

```
    Tested-by: John Garry <john.garry@huawei.com>
    Reviewed-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c446ec4ea8ae8f32b153e9917425c5dcd5b23a26
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Jun 2 21:30:28 2021 +0800

    scsi: core: Put .shost_dev in failure path if host state changes to RUNNING

    commit 11714026c02d613c30a149c3f4c4a15047744529 upstream.

    scsi_host_dev_release() only frees dev_name when host state is
    SHOST_CREATED. After host state has changed to SHOST_RUNNING,
    scsi_host_dev_release() no longer cleans up.

    Fix this by doing a put_device(&shost->shost_dev) in the failure path when
    host state is SHOST_RUNNING. Move get_device(&shost->shost_gendev) before
    device_add(&shost->shost_dev) so that scsi_host_cls_release() can do a put
    on this reference.

    Link: https://lore.kernel.org/r/20210602133029.2864069-4-ming.lei@redhat.com
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: Hannes Reinecke <hare@suse.de>
    Reported-by: John Garry <john.garry@huawei.com>
    Tested-by: John Garry <john.garry@huawei.com>
    Reviewed-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4aaa0d8efbefe7b10fdb2c0440c998227b15d58c
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Jun 2 21:30:27 2021 +0800

    scsi: core: Fix failure handling of scsi_add_host_with_dma()

    commit 3719f4ff047e20062b8314c23ec3cab84d74c908 upstream.

    When scsi_add_host_with_dma() returns failure, the caller will call
    scsi_host_put(shost) to release everything allocated for this host
    instance. Consequently we can't also free allocated stuff in
    scsi_add_host_with_dma(), otherwise we will end up with a double free.

    Strictly speaking, host resource allocations should have been done in
    scsi_host_alloc(). However, the allocations may need information which is
    not yet provided by the driver when that function is called. So leave the
    allocations where they are but rely on host device's release handler to
    free resources.

    Link: https://lore.kernel.org/r/20210602133029.2864069-3-ming.lei@redhat.com
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: John Garry <john.garry@huawei.com>
    Cc: Hannes Reinecke <hare@suse.de>
    Tested-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Reviewed-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 45d83db4728127944b237c0c8248987df9d478e7
Author: Ming Lei <ming.lei@redhat.com>
Date:   Wed Jun 2 21:30:26 2021 +0800

    scsi: core: Fix error handling of scsi_host_alloc()

    commit 66a834d092930cf41d809c0e989b13cd6f9ca006 upstream.

    After device is initialized via device_initialize(), or its name is set via
    dev_set_name(), the device has to be freed via put_device().  Otherwise
    device name will be leaked because it is allocated dynamically in
    dev_set_name().

    Fix the leak by replacing kfree() with put_device(). Since
    scsi_host_dev_release() properly handles IDA and kthread removal, remove
    special-casing these from the error handling as well.

    Link: https://lore.kernel.org/r/20210602133029.2864069-2-ming.lei@redhat.com
    Cc: Bart Van Assche <bvanassche@acm.org>
    Cc: John Garry <john.garry@huawei.com>
    Cc: Hannes Reinecke <hare@suse.de>
    Tested-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Reviewed-by: John Garry <john.garry@huawei.com>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Signed-off-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2431180f5c32775ad95cc875e094e7865e13ec90
Author: Dai Ngo <dai.ngo@oracle.com>
Date:   Wed May 19 17:15:10 2021 -0400

    NFSv4: nfs4_proc_set_acl needs to restore NFS_CAP_UIDGID_NOMAP on error.

    commit f8849e206ef52b584cd9227255f4724f0cc900bb upstream.

    Currently if __nfs4_proc_set_acl fails with NFS4ERR_BADOWNER it
    re-enables the idmapper by clearing NFS_CAP_UIDGID_NOMAP before
    retrying again. The NFS_CAP_UIDGID_NOMAP remains cleared even if
    the retry fails. This causes problem for subsequent setattr
    requests for v4 server that does not have idmapping configured.

    This patch modifies nfs4_proc_set_acl to detect NFS4ERR_BADOWNER
    and NFS4ERR_BADNAME and skips the retry, since the kernel isn't
    involved in encoding the ACEs, and return -EINVAL.

    Steps to reproduce the problem:

     # mount -o vers=4.1,sec=sys server:/export/test /tmp/mnt
     # touch /tmp/mnt/file1
     # chown 99 /tmp/mnt/file1
     # nfs4_setfacl -a A::unknown.user@xyz.com:wrtncy /tmp/mnt/file1
     Failed setxattr operation: Invalid argument
     # chown 99 /tmp/mnt/file1
     chown: changing ownership of ‘/tmp/mnt/file1’: Invalid argument
     # umount /tmp/mnt
     # mount -o vers=4.1,sec=sys server:/export/test /tmp/mnt
     # chown 99 /tmp/mnt/file1
     #

    v2: detect NFS4ERR_BADOWNER and NFS4ERR_BADNAME and skip retry
           in nfs4_proc_set_acl.
    Signed-off-by: Dai Ngo <dai.ngo@oracle.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
```

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

commit 64c21cd5c45d54f19b1bb4bf19d2c2d69f8f68e4
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Tue Jun 1 11:35:56 2021 -0400

    NFSv4: Fix second deadlock in nfs4_evict_inode()

    commit c3aba897c6e67fa464ec02b1f17911577d619713 upstream.

    If the inode is being evicted but has to return a layout first, then
    that too can cause a deadlock in the corner case where the server
    reboots.

    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 72651c6579a25317a90536181d311c663d0329ab
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Wed Jun 2 14:31:20 2021 -0400

    NFS: Fix use-after-free in nfs4_init_client()

    commit 476bdb04c501fc64bf3b8464ffddefc8dbe01577 upstream.

    KASAN reports a use-after-free when attempting to mount two different
    exports through two different NICs that belong to the same server.

    Olga was able to hit this with kernels starting somewhere between 5.7
    and 5.10, but I traced the patch that introduced the clear_bit() call to
    4.13. So something must have changed in the refcounting of the clp
    pointer to make this call to nfs_put_client() the very last one.

    Fixes: 8dcbec6d20 ("NFSv41: Handle EXCHID4_FLAG_CONFIRMED_R during NFSv4.1 migration")
    Cc: stable@vger.kernel.org # 4.13+
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 13f3c16b88eea3aa6890a4ca48843f716c724e86
Author: Paolo Bonzini <pbonzini@redhat.com>
Date:   Wed Jun 9 01:49:13 2021 -0400

    kvm: fix previous commit for 32-bit builds

    commit 4422829e8053068e0225e4d0ef42dc41ea7c9ef5 upstream.

    array_index_nospec does not work for uint64_t on 32-bit builds.
    However, the size of a memory slot must be less than 20 bits wide
    on those system, since the memory slot must fit in the user
    address space.  So just store it in an unsigned long.

    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e3f498b5d9f5bc1a9c3e301e46095c2c5e19e76
Author: Leo Yan <leo.yan@linaro.org>
Date:   Sat Jun 5 13:29:57 2021 +0800

    perf session: Correct buffer copying when peeking events

    [ Upstream commit 197eecb6ecae0b04bd694432f640ff75597fed9c ]

    When peeking an event, it has a short path and a long path.  The short
    path uses the session pointer "one_mmap_addr" to directly fetch the
```

```
    event; and the long path needs to read out the event header and the
    following event data from file and fill into the buffer pointer passed
    through the argument "buf".

    The issue is in the long path that it copies the event header and event
    data into the same destination address which pointer "buf", this means
    the event header is overwritten.  We are just lucky to run into the
    short path in most cases, so we don't hit the issue in the long path.

    This patch adds the offset "hdr_sz" to the pointer "buf" when copying
    the event data, so that it can reserve the event header which can be
    used properly by its caller.

    Fixes: 5a52f33adf02 ("perf session: Add perf_session__peek_event()")
    Signed-off-by: Leo Yan <leo.yan@linaro.org>
    Acked-by: Adrian Hunter <adrian.hunter@intel.com>
    Acked-by: Jiri Olsa <jolsa@redhat.com>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Kan Liang <kan.liang@linux.intel.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/20210605052957.1070720-1-leo.yan@linaro.org
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ff7b80b917c41d46d21ea285f42c2ed8626f2a7e
Author: Bjorn Andersson <bjorn.andersson@linaro.org>
Date:   Tue Jun 8 11:07:02 2021 -0700

    pinctrl: qcom: Make it possible to select SC8180x TLMM

    [ Upstream commit 30e9857a134905ac0d03ca244b615cc3ff0a076e ]

    It's currently not possible to select the SC8180x TLMM driver, due to it
    selecting PINCTRL_MSM, rather than depending on the same. Fix this.

    Fixes: 97423113ec4b ("pinctrl: qcom: Add sc8180x TLMM driver")
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20210608180702.2064253-1-bjorn.andersson@linaro.org
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc3406d4f7c056c4e65769052d42acdaa3be674a
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Tue Jun 1 11:10:05 2021 -0400

    NFSv4: Fix deadlock between nfs4_evict_inode() and nfs4_opendata_get_inode()

    [ Upstream commit dfe1fe75e00e4c724ede7b9e593f6f680e446c5f ]

    If the inode is being evicted, but has to return a delegation first,
    then it can cause a deadlock in the corner case where the server reboots
    before the delegreturn completes, but while the call to iget5_locked() in
    nfs4_opendata_get_inode() is waiting for the inode free to complete.
    Since the open call still holds a session slot, the reboot recovery
    cannot proceed.

    In order to break the logjam, we can turn the delegation return into a
    privileged operation for the case where we're evicting the inode. We
    know that in that case, there can be no other state recovery operation
    that conflicts.

    Reported-by: zhangxiaoxu (A) <zhangxiaoxu5@huawei.com>
    Fixes: 5fcdfacc01f3 ("NFSv4: Return delegations synchronously in evict_inode")
```

```
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 58ddf61f10b8f9b7b1341644bfee2f1c6508d4e1
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Thu Jun 3 15:37:53 2021 +0300

    NFS: Fix a potential NULL dereference in nfs_get_client()

    [ Upstream commit 09226e8303beeec10f2ff844d2e46d1371dc58e0 ]

    None of the callers are expecting NULL returns from nfs_get_client() so
    this code will lead to an Oops.  It's better to return an error
    pointer.  I expect that this is dead code so hopefully no one is
    affected.

    Fixes: 31434f496abb ("nfs: check hostname in nfs_get_client")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3e670c54eda238cb8a1ea93538a79ae89285c1c4
Author: Alaa Hleihel <alaa@nvidia.com>
Date:   Thu Jun 10 10:34:27 2021 +0300

    IB/mlx5: Fix initializing CQ fragments buffer

    commit 2ba0aa2feebda680ecfc3c552e867cf4d1b05a3a upstream.

    The function init_cq_frag_buf() can be called to initialize the current CQ
    fragments buffer cq->buf, or the temporary cq->resize_buf that is filled
    during CQ resize operation.

    However, the offending commit started to use function get_cqe() for
    getting the CQEs, the issue with this change is that get_cqe() always
    returns CQEs from cq->buf, which leads us to initialize the wrong buffer,
    and in case of enlarging the CQ we try to access elements beyond the size
    of the current cq->buf and eventually hit a kernel panic.

     [exception RIP: init_cq_frag_buf+103]
      [ffff9f799ddcbcd8] mlx5_ib_resize_cq at ffffffffc0835d60 [mlx5_ib]
      [ffff9f799ddcbdb0] ib_resize_cq at ffffffffc05270df [ib_core]
      [ffff9f799ddcbdc0] llt_rdma_setup_qp at ffffffffc0a6a712 [llt]
      [ffff9f799ddcbe10] llt_rdma_cc_event_action at ffffffffc0a6b411 [llt]
      [ffff9f799ddcbe98] llt_rdma_client_conn_thread at ffffffffc0a6bb75 [llt]
      [ffff9f799ddcbec8] kthread at ffffffffa66c5da1
      [ffff9f799ddcbf50] ret_from_fork_nospec_begin at ffffffffa6d95ddd

    Fix it by getting the needed CQE by calling mlx5_frag_buf_get_wqe() that
    takes the correct source buffer as a parameter.

    Fixes: 388ca8be0037 ("IB/mlx5: Implement fragmented completion queue (CQ)")
    Link:
https://lore.kernel.org/r/90a0e8c924093cfa50a482880ad7e7edb73dc19a.1623309971.git.leonro@nvi
dia.com
    Signed-off-by: Alaa Hleihel <alaa@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9fb088ce13bc3c59a51260207b487db3e556f275
Author: Sean Christopherson <seanjc@google.com>
Date:   Mon Jun 7 10:57:48 2021 -0700

    KVM: x86: Ensure liveliness of nested VM-Enter fail tracepoint message
```

```
    commit f31500b0d437a2464ca5972d8f5439e156b74960 upstream.

    Use the __string() machinery provided by the tracing subystem to make a
    copy of the string literals consumed by the "nested VM-Enter failed"
    tracepoint.  A complete copy is necessary to ensure that the tracepoint
    can't outlive the data/memory it consumes and deference stale memory.

    Because the tracepoint itself is defined by kvm, if kvm-intel and/or
    kvm-amd are built as modules, the memory holding the string literals
    defined by the vendor modules will be freed when the module is unloaded,
    whereas the tracepoint and its data in the ring buffer will live until
    kvm is unloaded (or "indefinitely" if kvm is built-in).

    This bug has existed since the tracepoint was added, but was recently
    exposed by a new check in tracing to detect exactly this type of bug.

      fmt: '%s%s
      ' current_buffer: ' vmx_dirty_log_t-140127  [003] ....  kvm_nested_vmenter_failed: '
      WARNING: CPU: 3 PID: 140134 at kernel/trace/trace.c:3759
trace_check_vprintf+0x3be/0x3e0
      CPU: 3 PID: 140134 Comm: less Not tainted 5.13.0-rc1-ce2e73ce600a-req #184
      Hardware name: ASUS Q87M-E/Q87M-E, BIOS 1102 03/03/2014
      RIP: 0010:trace_check_vprintf+0x3be/0x3e0
      Code: <0f> 0b 44 8b 4c 24 1c e9 a9 fe ff ff c6 44 02 ff 00 49 8b 97 b0 20
      RSP: 0018:ffffa895cc37bcb0 EFLAGS: 00010282
      RAX: 0000000000000000 RBX: ffffa895cc37bd08 RCX: 0000000000000027
      RDX: 0000000000000027 RSI: 00000000ffffdfff RDI: ffff9766cfad74f8
      RBP: ffffffffc0a041d4 R08: ffff9766cfad74f0 R09: ffffa895cc37bad8
      R10: 0000000000000001 R11: 0000000000000001 R12: ffffffffc0a041d4
      R13: ffffffffc0f4dba8 R14: 0000000000000000 R15: ffff976409f2c000
      FS:  00007f92fa200740(0000) GS:ffff9766cfac0000(0000) knlGS:0000000000000000
      CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
      CR2: 0000559bd11b0000 CR3: 000000019fbaa002 CR4: 0000000001726e0
      Call Trace:
       trace_event_printf+0x5e/0x80
       trace_raw_output_kvm_nested_vmenter_failed+0x3a/0x60 [kvm]
       print_trace_line+0x1dd/0x4e0
       s_show+0x45/0x150
       seq_read_iter+0x2d5/0x4c0
       seq_read+0x106/0x150
       vfs_read+0x98/0x180
       ksys_read+0x5f/0xe0
       do_syscall_64+0x40/0xb0
       entry_SYSCALL_64_after_hwframe+0x44/0xae

    Cc: Steven Rostedt <rostedt@goodmis.org>
    Fixes: 380e0055bc7e ("KVM: nVMX: trace nested VM-Enter failures detected by H/W")
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Reviewed-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Message-Id: <20210607175748.674002-1-seanjc@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 22cf65b6902b53c3db75b3f4307312be678772c6
Author: CodyYao-oc <CodyYao-oc@zhaoxin.com>
Date:   Mon Jun 7 10:53:35 2021 +0800

    x86/nmi_watchdog: Fix old-style NMI watchdog regression on old Intel CPUs

    commit a8383dfb2138742a1bb77b481ada047aededa2ba upstream.

    The following commit:

       3a4ac121c2ca ("x86/perf: Add hardware performance events support for Zhaoxin CPU.")
```

```
    Got the old-style NMI watchdog logic wrong and broke it for basically every
    Intel CPU where it was active. Which is only truly old CPUs, so few people noticed.

    On CPUs with perf events support we turn off the old-style NMI watchdog, so it
    was pretty pointless to add the logic for X86_VENDOR_ZHAOXIN to begin with ... :-/

    Anyway, the fix is to restore the old logic and add a 'break'.

    [ mingo: Wrote a new changelog. ]

    Fixes: 3a4ac121c2ca ("x86/perf: Add hardware performance events support for Zhaoxin
CPU.")
    Signed-off-by: CodyYao-oc <CodyYao-oc@zhaoxin.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lore.kernel.org/r/20210607025335.9643-1-CodyYao-oc@zhaoxin.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 63345c030e715c91d39e696243922a0453918fad
Author: Maximilian Luz <luzmaximilian@gmail.com>
Date:   Thu Jun 3 02:06:36 2021 +0200

    platform/surface: aggregator: Fix event disable function

    commit b430e1d65ef6eeee42c4e53028f8dfcc6abc728b upstream.

    Disabling events silently fails due to the wrong command ID being used.
    Instead of the command ID for the disable call, the command ID for the
    enable call was being used. This causes the disable call to enable the
    event instead. As the event is already enabled when we call this
    function, the EC silently drops this command and does nothing.

    Use the correct command ID for disabling the event to fix this.

    Fixes: c167b9c7e3d6 ("platform/surface: Add Surface Aggregator subsystem")
    Signed-off-by: Maximilian Luz <luzmaximilian@gmail.com>
    Link: https://lore.kernel.org/r/20210603000636.568846-1-luzmaximilian@gmail.com
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 322fd1cfe16903c17cf36f8770d7310151694065
Author: Dietmar Eggemann <dietmar.eggemann@arm.com>
Date:   Wed Jun 2 16:58:08 2021 +0200

    sched/fair: Fix util_est UTIL_AVG_UNCHANGED handling

    commit 68d7a190682aa4eb02db477328088ebad15acc83 upstream.

    The util_est internal UTIL_AVG_UNCHANGED flag which is used to prevent
    unnecessary util_est updates uses the LSB of util_est.enqueued. It is
    exposed via _task_util_est() (and task_util_est()).

    Commit 92a801e5d5b7 ("sched/fair: Mask UTIL_AVG_UNCHANGED usages")
    mentions that the LSB is lost for util_est resolution but
    find_energy_efficient_cpu() checks if task_util_est() returns 0 to
    return prev_cpu early.

    _task_util_est() returns the max value of util_est.ewma and
    util_est.enqueued or'ed w/ UTIL_AVG_UNCHANGED.
    So task_util_est() returning the max of task_util() and
    _task_util_est() will never return 0 under the default
    SCHED_FEAT(UTIL_EST, true).

    To fix this use the MSB of util_est.enqueued instead and keep the flag
    util_est internal, i.e. don't export it via _task_util_est().
```

```
    The maximal possible util_avg value for a task is 1024 so the MSB of
    'unsigned int util_est.enqueued' isn't used to store a util value.

    As a caveat the code behind the util_est_se trace point has to filter
    UTIL_AVG_UNCHANGED to see the real util_est.enqueued value which should
    be easy to do.

    This also fixes an issue report by Xuewen Yan that util_est_update()
    only used UTIL_AVG_UNCHANGED for the subtrahend of the equation:

      last_enqueued_diff = ue.enqueued - (task_util() | UTIL_AVG_UNCHANGED)

    Fixes: b89997aa88f0b sched/pelt: Fix task util_est update filtering
    Signed-off-by: Dietmar Eggemann <dietmar.eggemann@arm.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Xuewen Yan <xuewen.yan@unisoc.com>
    Reviewed-by: Vincent Donnefort <vincent.donnefort@arm.com>
    Reviewed-by: Vincent Guittot <vincent.guittot@linaro.org>
    Link: https://lore.kernel.org/r/20210602145808.1562603-1-dietmar.eggemann@arm.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 996c2417151ee0032708997ec735aa477a2cf4a8
Author: Kai Vehmanen <kai.vehmanen@linux.intel.com>
Date:   Fri May 28 17:43:30 2021 +0300

    ASoC: SOF: reset enabled_cores state at suspend

    commit b640e8a4bd24e17ce24a064d704aba14831651a8 upstream.

    The recent changes to use common code to power up/down DSP cores also
    removed the reset of the core state at suspend. It turns out this is
    still needed. When the firmware state is reset to
    SOF_FW_BOOT_NOT_STARTED, also enabled_cores should be reset, and
    existing DSP drivers depend on this.

    BugLink: https://github.com/thesofproject/linux/issues/2824
    Fixes: 42077f08b3 ("ASoC: SOF: update dsp core power status in common APIs")
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Link: https://lore.kernel.org/r/20210528144330.2551-1-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5e3fb56e430b425df1621779abc6216631cce2e3
Author: Vincent Guittot <vincent.guittot@linaro.org>
Date:   Thu May 27 14:29:16 2021 +0200

    sched/fair: Make sure to update tg contrib for blocked load

    commit 02da26ad5ed6ea8680e5d01f20661439611ed776 upstream.

    During the update of fair blocked load (__update_blocked_fair()), we
    update the contribution of the cfs in tg->load_avg if cfs_rq's pelt
    has decayed.  Nevertheless, the pelt values of a cfs_rq could have
    been recently updated while propagating the change of a child. In this
    case, cfs_rq's pelt will not decayed because it has already been
    updated and we don't update tg->load_avg.

    __update_blocked_fair
      ...
      for_each_leaf_cfs_rq_safe: child cfs_rq
        update cfs_rq_load_avg() for child cfs_rq
        ...
```

```
          update_load_avg(cfs_rq_of(se), se, 0)
            ...
            update cfs_rq_load_avg() for parent cfs_rq
                      -propagation of child's load makes parent cfs_rq->load_sum
                       becoming null
              -UPDATE_TG is not set so it doesn't update parent
                       cfs_rq->tg_load_avg_contrib
        ..
        for_each_leaf_cfs_rq_safe: parent cfs_rq
          update cfs_rq_load_avg() for parent cfs_rq
            - nothing to do because parent cfs_rq has already been updated
                      recently so cfs_rq->tg_load_avg_contrib is not updated
          ...
          parent cfs_rq is decayed
            list_del_leaf_cfs_rq parent cfs_rq
                - but it still contibutes to tg->load_avg

    we must set UPDATE_TG flags when propagting pending load to the parent

    Fixes: 039ae8bcf7a5 ("sched/fair: Fix O(nr_cgroups) in the load balancing path")
    Reported-by: Odin Ugedal <odin@uged.al>
    Signed-off-by: Vincent Guittot <vincent.guittot@linaro.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Odin Ugedal <odin@uged.al>
    Link: https://lkml.kernel.org/r/20210527122916.27683-3-vincent.guittot@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2481ecc6835a4a7729ed0a33150bf46bd3242349
Author: Vincent Guittot <vincent.guittot@linaro.org>
Date:   Thu May 27 14:29:15 2021 +0200

    sched/fair: Keep load_avg and load_sum synced

    commit 7c7ad626d9a0ff0a36c1e2a3cfbbc6a13828d5eb upstream.

    when removing a cfs_rq from the list we only check _sum value so we must
    ensure that _avg and _sum stay synced so load_sum can't be null whereas
    load_avg is not after propagating load in the cgroup hierarchy.

    Use load_avg to compute load_sum similarly to what is done for util_sum
    and runnable_sum.

    Fixes: 0e2d2aaaae52 ("sched/fair: Rewrite PELT migration propagation")
    Reported-by: Odin Ugedal <odin@uged.al>
    Signed-off-by: Vincent Guittot <vincent.guittot@linaro.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Odin Ugedal <odin@uged.al>
    Link: https://lkml.kernel.org/r/20210527122916.27683-2-vincent.guittot@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 20d396b17a405b1df872822322a4cf759e9fedce
Author: Kyle Tso <kyletso@google.com>
Date:   Thu May 27 20:10:29 2021 +0800

    dt-bindings: connector: Replace BIT macro with generic bit ops

    commit 9257bd80b917cc7908abd27ed5a5211964563f62 upstream.

    BIT macro is not defined. Replace it with generic bit operations.

    Fixes: 630dce2810b9 ("dt-bindings: connector: Add SVDM VDO properties")
    Reviewed-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210527121029.583611-1-kyletso@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit df4c7e3b10957369b2180a49c25fdff73bee139b
Author: Marco Elver <elver@google.com>
Date:   Thu May 27 12:47:11 2021 +0200

    perf: Fix data race between pin_count increment/decrement

    commit 6c605f8371159432ec61cbb1488dcf7ad24ad19a upstream.

    KCSAN reports a data race between increment and decrement of pin_count:

      write to 0xffff888237c2d4e0 of 4 bytes by task 15740 on cpu 1:
       find_get_context             kernel/events/core.c:4617
       __do_sys_perf_event_open     kernel/events/core.c:12097 [inline]
       __se_sys_perf_event_open     kernel/events/core.c:11933
       ...
      read to 0xffff888237c2d4e0 of 4 bytes by task 15743 on cpu 0:
       perf_unpin_context           kernel/events/core.c:1525 [inline]
       __do_sys_perf_event_open     kernel/events/core.c:12328 [inline]
       __se_sys_perf_event_open     kernel/events/core.c:11933
       ...

    Because neither read-modify-write here is atomic, this can lead to one
    of the operations being lost, resulting in an inconsistent pin_count.
    Fix it by adding the missing locking in the CPU-event case.

    Fixes: fe4b04fa31a6 ("perf: Cure task_oncpu_function_call() races")
    Reported-by: syzbot+142c9018f5962db69c7e@syzkaller.appspotmail.com
    Signed-off-by: Marco Elver <elver@google.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lkml.kernel.org/r/20210527104711.2671610-1-elver@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5138f8df070011468105dd06b975ad18088d0de3
Author: Kan Liang <kan.liang@linux.intel.com>
Date:   Wed May 26 06:58:47 2021 -0700

    perf/x86/intel/uncore: Fix a kernel WARNING triggered by maxcpus=1

    commit 4a0e3ff30980b7601b13dd3b7ee275212b852843 upstream.

    A kernel WARNING may be triggered when setting maxcpus=1.

    The uncore counters are Die-scope. When probing a PCI device, only the
    BUS information can be retrieved. The uncore driver has to maintain a
    mapping table used to calculate the logical Die ID from a given BUS#.

    Before the patch ba9506be4e40, the mapping table stores the mapping
    information from the BUS# -> a Physical Socket ID. To calculate the
    logical die ID, perf does,
    - In snbep_pci2phy_map_init(), retrieve the BUS# -> a Physical Socket ID
      from the UBOX PCI configure space.
    - Calculate the mapping information (a BUS# -> a Physical Socket ID) for
      the other PCI BUS.
    - In the uncore_pci_probe(), get the physical Socket ID from a given BUS
      and the mapping table.
    - Calculate the logical Die ID

    Since only the logical Die ID is required, with the patch ba9506be4e40,
    the mapping table stores the mapping information from the BUS# -> a
    logical Die ID. Now perf does,
    - In snbep_pci2phy_map_init(), retrieve the BUS# -> a Physical Socket ID
      from the UBOX PCI configure space.
    - Calculate the logical Die ID
    - Calculate the mapping information (a BUS# -> a logical Die ID) for the
```

```
      other PCI BUS.
    - In the uncore_pci_probe(), get the logical die ID from a given BUS and
      the mapping table.

    When calculating the logical Die ID, -1 may be returned, especially when
    maxcpus=1. Here, -1 means the logical Die ID is not found. But when
    calculating the mapping information for the other PCI BUS, -1 indicates
    that it's the other PCI BUS that requires the calculation of the
    mapping. The driver will mistakenly do the calculation.

    Uses the -ENODEV to indicate the case which the logical Die ID is not
    found. The driver will not mess up the mapping table anymore.

    Fixes: ba9506be4e40 ("perf/x86/intel/uncore: Store the logical die id instead of the
physical die id.")
    Reported-by: John Donnelly <john.p.donnelly@oracle.com>
    Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Acked-by: John Donnelly <john.p.donnelly@oracle.com>
    Tested-by: John Donnelly <john.p.donnelly@oracle.com>
    Link: https://lkml.kernel.org/r/1622037527-156028-1-git-send-email-
kan.liang@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f6025f3ff1fdad4572d15648423c97c06dac43b
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed May 26 13:58:57 2021 +0530

    pinctrl: qcom: Fix duplication in gpio_groups

    commit 0e4bf265b11a00bde9fef6b791bd8ee2d8059701 upstream.

    "gpio52" and "gpio53" are duplicated in gpio_groups, fix them!

    Fixes: ac43c44a7a37 ("pinctrl: qcom: Add SDX55 pincontrol driver")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20210526082857.174682-1-manivannan.sadhasivam@linaro.org
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dd55331d493b7ea75c5db1f24d6822946fde2862
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Tue May 25 17:55:39 2021 +0100

    gpio: wcd934x: Fix shift-out-of-bounds error

    commit dbec64b11c65d74f31427e2b9d5746fbf17bf840 upstream.

    bit-mask for pins 0 to 4 is BIT(0) to BIT(4) however we ended up with BIT(n - 1)
    which is not right, and this was caught by below usban check

    UBSAN: shift-out-of-bounds in drivers/gpio/gpio-wcd934x.c:34:14

    Fixes: 59c324683400 ("gpio: wcd934x: Add support to wcd934x gpio controller")
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Bartosz Golaszewski <bgolaszewski@baylibre.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7307a2d3fdf8dc0ee35654f1c9e5cb6f3f72eb5a
Author: Yang Li <yang.lee@linux.alibaba.com>
Date:   Tue May 25 18:50:32 2021 +0800
```

```
    phy: ti: Fix an error code in wiz_probe()

    commit b8203ec7f58ae925e10fadd3d136073ae7503a6e upstream.

    When the code execute this if statement, the value of ret is 0.
    However, we can see from the dev_err() log that the value of
    ret should be -EINVAL.

    Clean up smatch warning:

    drivers/phy/ti/phy-j721e-wiz.c:1216 wiz_probe() warn: missing error code 'ret'

    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Fixes: c9f9eba06629 ("phy: ti: j721e-wiz: Manage typec-gpio-dir")
    Signed-off-by: Yang Li <yang.lee@linux.alibaba.com>
    Link: https://lore.kernel.org/r/1621939832-65535-1-git-send-email-
yang.lee@linux.alibaba.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 37d144cd4456f3318f2ccf5c7138f79e0fff2e2a
Author: Jerome Brunet <jbrunet@baylibre.com>
Date:   Mon May 24 11:34:48 2021 +0200

    ASoC: meson: gx-card: fix sound-dai dt schema

    commit d031d99b02eaf7363c33f5b27b38086cc8104082 upstream.

    There is a fair amount of warnings when running 'make dtbs_check' with
    amlogic,gx-sound-card.yaml.

    Ex:
    arch/arm64/boot/dts/amlogic/meson-gxm-q200.dt.yaml: sound: dai-link-0:sound-dai:0:1:
missing phandle tag in 0
    arch/arm64/boot/dts/amlogic/meson-gxm-q200.dt.yaml: sound: dai-link-0:sound-dai:0:2:
missing phandle tag in 0
    arch/arm64/boot/dts/amlogic/meson-gxm-q200.dt.yaml: sound: dai-link-0:sound-dai:0: [66,
0, 0] is too long

    The reason is that the sound-dai phandle provided has cells, and in such
    case the schema should use 'phandle-array' instead of 'phandle'.

    Fixes: fd00366b8e41 ("ASoC: meson: gx: add sound card dt-binding documentation")
    Signed-off-by: Jerome Brunet <jbrunet@baylibre.com>
    Link: https://lore.kernel.org/r/20210524093448.357140-1-jbrunet@baylibre.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 047fd16015a79180771650aa6ce71f68b2c23368
Author: Kefeng Wang <wangkefeng.wang@huawei.com>
Date:   Mon May 24 10:49:41 2021 +0800

    ASoC: core: Fix Null-point-dereference in fmt_single_name()

    commit 41daf6ba594d55f201c50280ebcd430590441da1 upstream.

    Check the return value of devm_kstrdup() in case of
    Null-point-dereference.

    Fixes: 45dd9943fce0 ("ASoC: core: remove artificial component and DAI name constraint")
    Cc: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Kefeng Wang <wangkefeng.wang@huawei.com>
    Link: https://lore.kernel.org/r/20210524024941.159952-1-wangkefeng.wang@huawei.com
```

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

commit 02b07dd201f0322ecd5475715f92de24b085df14
Author: Wang Wensheng <wangwensheng4@huawei.com>
Date:   Mon May 17 01:57:49 2021 +0000

    phy: cadence: Sierra: Fix error return code in cdns_sierra_phy_probe()

    commit 6411e386db0a477217607015e7d2910d02f75426 upstream.

    Fix to return a negative error code from the error handling
    case instead of 0, as done elsewhere in this function.

    Fixes: a43f72ae136a ("phy: cadence: Sierra: Change MAX_LANES of Sierra to 16")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wang Wensheng <wangwensheng4@huawei.com>
    Link: https://lore.kernel.org/r/20210517015749.127799-1-wangwensheng4@huawei.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c8783017306a7f460c1e74f81aee0d367b4f80c
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Sat May 8 11:42:16 2021 +0800

    tools/bootconfig: Fix error return code in apply_xbc()

    commit e8ba0b2b64126381643bb50df3556b139a60545a upstream.

    Fix to return a negative error code from the error handling case instead
    of 0, as done elsewhere in this function.

    Link: https://lkml.kernel.org/r/20210508034216.2277-1-thunder.leizhen@huawei.com

    Fixes: a995e6bc0524 ("tools/bootconfig: Fix to check the write failure correctly")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 05471eb76e777bf70bc52547fe732e7f2337ae3b
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed May 5 17:14:11 2021 -0700

    vmlinux.lds.h: Avoid orphan section with !SMP

    commit d4c6399900364facd84c9e35ce1540b6046c345f upstream.

    With x86_64_defconfig and the following configs, there is an orphan
    section warning:

    CONFIG_SMP=n
    CONFIG_AMD_MEM_ENCRYPT=y
    CONFIG_HYPERVISOR_GUEST=y
    CONFIG_KVM=y
    CONFIG_PARAVIRT=y

    ld: warning: orphan section `.data..decrypted' from `arch/x86/kernel/cpu/vmware.o' being
placed in section `.data..decrypted'
    ld: warning: orphan section `.data..decrypted' from `arch/x86/kernel/kvm.o' being placed
in section `.data..decrypted'

    These sections are created with DEFINE_PER_CPU_DECRYPTED, which
    ultimately turns into __PCPU_ATTRS, which in turn has a section
```

```
    attribute with a value of PER_CPU_BASE_SECTION + the section name. When
    CONFIG_SMP is not set, the base section is .data and that is not
    currently handled in any linker script.

    Add .data..decrypted to PERCPU_DECRYPTED_SECTION, which is included in
    PERCPU_INPUT -> PERCPU_SECTION, which is include in the x86 linker
    script when either CONFIG_X86_64 or CONFIG_SMP is unset, taking care of
    the warning.

    Fixes: ac26963a1175 ("percpu: Introduce DEFINE_PER_CPU_DECRYPTED")
    Link: https://github.com/ClangBuiltLinux/linux/issues/1360
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Tested-by: Nick Desaulniers <ndesaulniers@google.com> # build
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210506001410.1026691-1-nathan@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d37fa2a044d3497ccb302a657d38b9f390f09996
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Wed Dec 30 16:54:56 2020 +0100

    ARM: cpuidle: Avoid orphan section warning

    commit d94b93a9101573eb75b819dee94b1417acff631b upstream.

    Since commit 83109d5d5fba ("x86/build: Warn on orphan section placement"),
    we get a warning for objects in orphan sections. The cpuidle implementation
    for OMAP causes this when CONFIG_CPU_IDLE is disabled:

    arm-linux-gnueabi-ld: warning: orphan section `__cpuidle_method_of_table' from
`arch/arm/mach-omap2/pm33xx-core.o' being placed in section `__cpuidle_method_of_table'
    arm-linux-gnueabi-ld: warning: orphan section `__cpuidle_method_of_table' from
`arch/arm/mach-omap2/pm33xx-core.o' being placed in section `__cpuidle_method_of_table'
    arm-linux-gnueabi-ld: warning: orphan section `__cpuidle_method_of_table' from
`arch/arm/mach-omap2/pm33xx-core.o' being placed in section `__cpuidle_method_of_table'

    Change the definition of CPUIDLE_METHOD_OF_DECLARE() to silently
    drop the table and all code referenced from it when CONFIG_CPU_IDLE
    is disabled.

    Fixes: 06ee7a950b6a ("ARM: OMAP2+: pm33xx-core: Add cpuidle_ops for am335x/am437x")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Miguel Ojeda <ojeda@kernel.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20201230155506.1085689-1-arnd@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8dc1b0e0ca204596c50bcd159ee069ae0f998176
Author: Maor Gottlieb <maorg@nvidia.com>
Date:   Thu Jun 10 10:34:25 2021 +0300

    RDMA: Verify port when creating flow rule

    commit 2adcb4c5a52a2623cd2b43efa7041e74d19f3a5e upstream.

    Validate port value provided by the user and with that remove no longer
    needed validation by the driver.  The missing check in the mlx5_ib driver
    could cause to the below oops.

    Call trace:
      _create_flow_rule+0x2d4/0xf28 [mlx5_ib]
      mlx5_ib_create_flow+0x2d0/0x5b0 [mlx5_ib]
      ib_uverbs_ex_create_flow+0x4cc/0x624 [ib_uverbs]
```

```
       ib_uverbs_handler_UVERBS_METHOD_INVOKE_WRITE+0xd4/0x150 [ib_uverbs]
       ib_uverbs_cmd_verbs.isra.7+0xb28/0xc50 [ib_uverbs]
       ib_uverbs_ioctl+0x158/0x1d0 [ib_uverbs]
       do_vfs_ioctl+0xd0/0xaf0
       ksys_ioctl+0x84/0xb4
       __arm64_sys_ioctl+0x28/0xc4
       el0_svc_common.constprop.3+0xa4/0x254
       el0_svc_handler+0x84/0xa0
       el0_svc+0x10/0x26c
      Code: b9401260 f9615681 51000400 8b001c20 (f9403c1a)

    Fixes: 436f2ad05a0b ("IB/core: Export ib_create/destroy_flow through uverbs")
    Link:
https://lore.kernel.org/r/faad30dc5219a01727f47db3dc2f029d07c82c00.1623309971.git.leonro@nvi
dia.com
    Reviewed-by: Mark Bloch <markb@mellanox.com>
    Signed-off-by: Maor Gottlieb <maorg@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fdeb9f8043e14abbbf0c22f7e8054ebcc5aa3d02
Author: Mark Bloch <mbloch@nvidia.com>
Date:   Mon Jun 7 11:03:12 2021 +0300

    RDMA/mlx5: Block FDB rules when not in switchdev mode

    commit edc0b0bccc9c80d9a44d3002dcca94984b25e7cf upstream.

    Allow creating FDB steering rules only when in switchdev mode.

    The only software model where a userspace application can manipulate
    FDB entries is when it manages the eswitch. This is only possible in
    switchdev mode where we expose a single RDMA device with representors
    for all the vports that are connected to the eswitch.

    Fixes: 52438be44112 ("RDMA/mlx5: Allow inserting a steering rule to the FDB")
    Link:
https://lore.kernel.org/r/e928ae7c58d07f104716a2a8d730963d1bd01204.1623052923.git.leonro@nvi
dia.com
    Reviewed-by: Maor Gottlieb <maorg@nvidia.com>
    Signed-off-by: Mark Bloch <mbloch@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 73b27339cb588d16fad43061321f20c9f1e095b3
Author: Shay Drory <shayd@nvidia.com>
Date:   Thu Jun 3 16:19:39 2021 +0300

    RDMA/mlx4: Do not map the core_clock page to user space unless enabled

    commit 404e5a12691fe797486475fe28cc0b80cb8bef2c upstream.

    Currently when mlx4 maps the hca_core_clock page to the user space there
    are read-modifiable registers, one of which is semaphore, on this page as
    well as the clock counter. If user reads the wrong offset, it can modify
    the semaphore and hang the device.

    Do not map the hca_core_clock page to the user space unless the device has
    been put in a backwards compatibility mode to support this feature.

    After this patch, mlx4 core_clock won't be mapped to user space on the
    majority of existing devices and the uverbs device time feature in
    ibv_query_rt_values_ex() will be disabled.
```

```
    Fixes: 52033cfb5aab ("IB/mlx4: Add mmap call to map the hardware clock")
    Link:
https://lore.kernel.org/r/9632304e0d6790af84b3b706d8c18732bc0d5e27.1622726305.git.leonro@nvi
dia.com
    Signed-off-by: Shay Drory <shayd@nvidia.com>
    Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 98b6c670ac2ddd1cdde2460c3210c686d04422f3
Author: Mark Zhang <markzhang@nvidia.com>
Date:   Thu Jun 3 16:18:03 2021 +0300

    RDMA/mlx5: Use different doorbell memory for different processes

    commit a0ffb4c12f7fa89163e228e6f27df09b46631db1 upstream.

    In a fork scenario, the parent and child can have same virtual address and
    also share the uverbs fd.  That causes to the list_for_each_entry search
    return same doorbell physical page for all processes, even though that
    page has been COW' or copied.

    This patch takes the mm_struct into consideration during search, to make
    sure that VA's belonging to different processes are not intermixed.

    Resolves the malfunction of uverbs after fork in some specific cases.

    Fixes: e126ba97dba9 ("mlx5: Add driver for Mellanox Connect-IB adapters")
    Link:
https://lore.kernel.org/r/feacc23fe0bc6e1088c6824d5583798745e72405.1622726212.git.leonro@nvi
dia.com
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-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: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0a672f7d89db2da17ae02733ccc08458be72a6f8
Author: Kamal Heib <kamalheib1@gmail.com>
Date:   Tue May 25 18:01:34 2021 +0300

    RDMA/ipoib: Fix warning caused by destroying non-initial netns

    commit a3e74fb9247cd530dca246699d5eb5a691884d32 upstream.

    After the commit 5ce2dced8e95 ("RDMA/ipoib: Set rtnl_link_ops for ipoib
    interfaces"), if the IPoIB device is moved to non-initial netns,
    destroying that netns lets the device vanish instead of moving it back to
    the initial netns, This is happening because default_device_exit() skips
    the interfaces due to having rtnl_link_ops set.

    Steps to reporoduce:
      ip netns add foo
      ip link set mlx5_ib0 netns foo
      ip netns delete foo

    WARNING: CPU: 1 PID: 704 at net/core/dev.c:11435 netdev_exit+0x3f/0x50
    Modules linked in: xt_CHECKSUM xt_MASQUERADE xt_conntrack ipt_REJECT
    nf_reject_ipv4 nft_compat nft_counter nft_chain_nat nf_nat nf_conntrack
    nf_defrag_ipv6 nf_defrag_ipv4 nf_tables nfnetlink tun d
     fuse
    CPU: 1 PID: 704 Comm: kworker/u64:3 Tainted: G S      W  5.13.0-rc1+ #1
    Hardware name: Dell Inc. PowerEdge R630/02C2CP, BIOS 2.1.5 04/11/2016
    Workqueue: netns cleanup_net
```

```
    RIP: 0010:netdev_exit+0x3f/0x50
    Code: 48 8b bb 30 01 00 00 e8 ef 81 b1 ff 48 81 fb c0 3a 54 a1 74 13 48
    8b 83 90 00 00 00 48 81 c3 90 00 00 00 48 39 d8 75 02 5b c3 <0f> 0b 5b
    c3 66 66 2e 0f 1f 84 00 00 00 00 00 66 90 0f 1f 44 00
    RSP: 0018:ffffb297079d7e08 EFLAGS: 00010206
    RAX: ffff8eb542c00040 RBX: ffff8eb541333150 RCX: 000000008010000d
    RDX: 000000008010000e RSI: 000000008010000d RDI: ffff8eb440042c00
    RBP: ffffb297079d7e48 R08: 0000000000000001 R09: ffffffff9fdeac00
    R10: ffff8eb5003be000 R11: 0000000000000001 R12: ffffffffa1545620
    R13: ffffffffa1545628 R14: 0000000000000000 R15: ffffffffa1543b20
    FS:  0000000000000000(0000) GS:ffff8ed37fa00000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00005601b5f4c2e8 CR3: 0000001fc8c10002 CR4: 00000000003706e0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     ops_exit_list.isra.9+0x36/0x70
     cleanup_net+0x234/0x390
     process_one_work+0x1cb/0x360
     ? process_one_work+0x360/0x360
     worker_thread+0x30/0x370
     ? process_one_work+0x360/0x360
     kthread+0x116/0x130
     ? kthread_park+0x80/0x80
     ret_from_fork+0x22/0x30

    To avoid the above warning and later on the kernel panic that could happen
    on shutdown due to a NULL pointer dereference, make sure to set the
    netns_refund flag that was introduced by commit 3a5ca857079e ("can: dev:
    Move device back to init netns on owning netns delete") to properly
    restore the IPoIB interfaces to the initial netns.

    Fixes: 5ce2dced8e95 ("RDMA/ipoib: Set rtnl_link_ops for ipoib interfaces")
    Link: https://lore.kernel.org/r/20210525150134.139342-1-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: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1ed3c16fe6b5ba1f2c37df602eebecdb7afc8e86
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Thu Jun 10 00:07:28 2021 +0200

    hwmon: (tps23861) correct shunt LSB values

    commit e13d1127241404f1c3eb1379ac4dd100eaf385b4 upstream.

    Current shunt LSB values got reversed during in the
    original driver commit.

    So, correct the current shunt LSB values according to
    the datasheet.

    This caused reading slightly skewed current values.

    Fixes: fff7b8ab2255 ("hwmon: add Texas Instruments TPS23861 driver")
    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Link: https://lore.kernel.org/r/20210609220728.499879-3-robert.marko@sartura.hr
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7ee186a7ead16ade2eee990c646b3ebb6ca47e32
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Thu Jun 10 00:07:27 2021 +0200
```

```
    hwmon: (tps23861) set current shunt value

    commit b325d3526e14942d42c392c2ac9fbea59c22894c upstream.

    TPS23861 has a configuration bit for setting of the
    current shunt value used on the board.
    Its bit 0 of the General Mask 1 register.

    According to the datasheet bit values are:
    0 for 255 mOhm (Default)
    1 for 250 mOhm

    So, configure the bit before registering the hwmon
    device according to the value passed in the DTS or
    default one if none is passed.

    This caused potentially reading slightly skewed values
    due to max current value being 1.02A when 250mOhm shunt
    is used instead of 1.0A when 255mOhm is used.

    Fixes: fff7b8ab2255 ("hwmon: add Texas Instruments TPS23861 driver")
    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Link: https://lore.kernel.org/r/20210609220728.499879-2-robert.marko@sartura.hr
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c1723ce28c5def05cec483698188fe4d26548856
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Thu Jun 10 00:07:26 2021 +0200

    hwmon: (tps23861) define regmap max register

    commit fb8543fb863e89baa433b4d716d73395caa1b7f4 upstream.

    Define the max register address the device supports.
    This allows reading the whole register space via
    regmap debugfs, without it only register 0x0 is visible.

    This was forgotten in the original driver commit.

    Fixes: fff7b8ab2255 ("hwmon: add Texas Instruments TPS23861 driver")
    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Link: https://lore.kernel.org/r/20210609220728.499879-1-robert.marko@sartura.hr
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e3b478db63dc64f40f231e862c68a0418a2378d6
Author: Wilken Gottwalt <wilken.gottwalt@posteo.net>
Date:   Thu Jun 3 11:51:02 2021 +0000

    hwmon: (corsair-psu) fix suspend behavior

    commit 7656cd2177612aa7c299b083ecff30a4d3e9a587 upstream.

    During standby some PSUs turn off the microcontroller. A re-init is
    required during resume or the microcontroller stays unresponsive.

    Fixes: d115b51e0e56 ("hwmon: add Corsair PSU HID controller driver")
    Signed-off-by: Wilken Gottwalt <wilken.gottwalt@posteo.net>
    Link: https://lore.kernel.org/r/YLjCJiVtu5zgTabI@monster.powergraphx.local
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0c010eae2febe81d7d25d554471cc015b950e0a1
Author: Jonathan Marek <jonathan@marek.ca>
```

```
Date:   Thu May 13 13:14:00 2021 -0400

    drm/msm/a6xx: avoid shadow NULL reference in failure path

    commit ce86c239e4d218ae6040bec18e6d19a58edb8b7c upstream.

    If a6xx_hw_init() fails before creating the shadow_bo, the a6xx_pm_suspend
    code referencing it will crash. Change the condition to one that avoids
    this problem (note: creation of shadow_bo is behind this same condition)

    Fixes: e8b0b994c3a5 ("drm/msm/a6xx: Clear shadow on suspend")
    Signed-off-by: Jonathan Marek <jonathan@marek.ca>
    Reviewed-by: Akhil P Oommen <akhilpo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210513171431.18632-6-jonathan@marek.ca
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 868fb39fbfce42c0c56779c137eb89442262a1f6
Author: Jonathan Marek <jonathan@marek.ca>
Date:   Thu May 13 13:13:59 2021 -0400

    drm/msm/a6xx: update/fix CP_PROTECT initialization

    commit 408434036958699a7f50ddec984f7ba33e11a8f5 upstream.

    Update CP_PROTECT register programming based on downstream.

    A6XX_PROTECT_RW is renamed to A6XX_PROTECT_NORDWR to make things aligned
    and also be more clear about what it does.

    Note that this required switching to use the CP_ALWAYS_ON_COUNTER as the
    GMU counter is not accessible from the cmdstream.  Which also means
    using the CPU counter for the msm_gpu_submit_flush() tracepoint (as
    catapult depends on being able to compare this to the start/end values
    captured in cmdstream).  This may need to be revisited when IFPC is
    enabled.

    Also, compared to downstream, this opens up CP_PERFCTR_CP_SEL as the
    userspace performance tooling (fdperf and pps-producer) expect to be
    able to configure the CP counters.

    Fixes: 4b565ca5a2cb ("drm/msm: Add A6XX device support")
    Signed-off-by: Jonathan Marek <jonathan@marek.ca>
    Reviewed-by: Akhil P Oommen <akhilpo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210513171431.18632-5-jonathan@marek.ca
    [switch to CP_ALWAYS_ON_COUNTER, open up CP_PERFCNTR_CP_SEL, and spiff
     up commit msg]
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 107107372f750482be4db57ef7c72b1cf5213950
Author: Jonathan Marek <jonathan@marek.ca>
Date:   Thu May 13 13:13:58 2021 -0400

    drm/msm/a6xx: fix incorrectly set uavflagprd_inv field for A650

    commit b4387eaf3821a4c4241ac3a556e13244eb1fdaa5 upstream.

    Value was shifted in the wrong direction, resulting in the field always
    being zero, which is incorrect for A650.

    Fixes: d0bac4e9cd66 ("drm/msm/a6xx: set ubwc config for A640 and A650")
    Signed-off-by: Jonathan Marek <jonathan@marek.ca>
    Reviewed-by: Akhil P Oommen <akhilpo@codeaurora.org>
    Link: https://lore.kernel.org/r/20210513171431.18632-4-jonathan@marek.ca
```

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

commit b6ed98736fdf5f14bb3ebbe2e676b976612291cd
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Tue Jun 8 23:33:18 2021 +0200

    drm/mcde: Fix off by 10^3 in calculation

    commit c8a570443943304cac2e4186dbce6989b6c2b8b5 upstream.

    The calclulation of how many bytes we stuff into the
    DSI pipeline for video mode panels is off by three
    orders of magnitude because we did not account for the
    fact that the DRM mode clock is in kilohertz rather
    than hertz.

    This used to be:
    drm_mode_vrefresh(mode) * mode->htotal * mode->vtotal
    which would become for example for s6e63m0:
    60 x 514 x 831 = 25628040 Hz, but mode->clock is
    25628 as it is in kHz.

    This affects only the Samsung GT-I8190 "Golden" phone
    right now since it is the only MCDE device with a video
    mode display.

    Curiously some specimen work with this code and wild
    settings in the EOL and empty packets at the end of the
    display, but I have noticed an eeire flicker until now.
    Others were not so lucky and got black screens.

    Cc: Ville Syrjälä <ville.syrjala@linux.intel.com>
    Reported-by: Stephan Gerhold <stephan@gerhold.net>
    Fixes: 920dd1b1425b ("drm/mcde: Use mode->clock instead of reverse calculating it from
the vrefresh")
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Tested-by: Stephan Gerhold <stephan@gerhold.net>
    Reviewed-by: Stephan Gerhold <stephan@gerhold.net>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210608213318.3897858-1-
linus.walleij@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 95aadff4afa00b3c362dafe0b7fb948c513d0d7d
Author: Bjorn Andersson <bjorn.andersson@linaro.org>
Date:   Wed Jun 9 17:21:32 2021 -0700

    usb: typec: mux: Fix copy-paste mistake in typec_mux_match

    commit 142d0b24c1b17139f1aaaacae7542a38aa85640f upstream.

    Fix the copy-paste mistake in the return path of typec_mux_match(),
    where dev is considered a member of struct typec_switch rather than
    struct typec_mux.

    The two structs are identical in regards to having the struct device as
    the first entry, so this provides no functional change.

    Fixes: 3370db35193b ("usb: typec: Registering real device entries for the muxes")
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20210610002132.3088083-1-bjorn.andersson@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0908fb0e80d8385c2ec3b3e94b136d85fdbfeee9
```

```
Author: Kyle Tso <kyletso@google.com>
Date:   Tue Jun 1 20:31:51 2021 +0800

    usb: typec: tcpm: Fix misuses of AMS invocation

    commit 80137c18737c30d20ee630e442405236d96898a7 upstream.

    tcpm_ams_start is used to initiate an AMS as well as checking Collision
    Avoidance conditions but not for flagging passive AMS (initiated by the
    port partner). Fix the misuses of tcpm_ams_start in tcpm_pd_svdm.

    ATTENTION doesn't need responses so the AMS flag is not needed here.

    Fixes: 0bc3ee92880d ("usb: typec: tcpm: Properly interrupt VDM AMS")
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210601123151.3441914-5-kyletso@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit aad940c6579cae757c11e0f306ea296e4ba6fcb3
Author: Kyle Tso <kyletso@google.com>
Date:   Tue Jun 1 20:31:48 2021 +0800

    usb: typec: tcpm: Correct the responses in SVDM Version 2.0 DFP

    commit f41bfc7e9c7c1d721c8752f1853cde43e606ad43 upstream.

    In USB PD Spec Rev 3.1 Ver 1.0, section "6.12.5 Applicability of
    Structured VDM Commands", DFP is allowed and recommended to respond to
    Discovery Identity with ACK. And in section "6.4.4.2.5.1 Commands other
    than Attention", NAK should be returned only when receiving Messages
    with invalid fields, Messages in wrong situation, or unrecognize
    Messages.

    Still keep the original design for SVDM Version 1.0 for backward
    compatibilities.

    Fixes: 193a68011fdc ("staging: typec: tcpm: Respond to Discover Identity commands")
    Acked-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210601123151.3441914-2-kyletso@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 63784c42dd3d0e5e90a64422480ad0b559be20b8
Author: Wesley Cheng <wcheng@codeaurora.org>
Date:   Thu May 20 21:23:57 2021 -0700

    usb: dwc3: gadget: Disable gadget IRQ during pullup disable

    commit 8212937305f84ef73ea81036dafb80c557583d4b upstream.

    Current sequence utilizes dwc3_gadget_disable_irq() alongside
    synchronize_irq() to ensure that no further DWC3 events are generated.
    However, the dwc3_gadget_disable_irq() API only disables device
    specific events.  Endpoint events can still be generated.  Briefly
    disable the interrupt line, so that the cleanup code can run to
    prevent device and endpoint events. (i.e. __dwc3_gadget_stop() and
    dwc3_stop_active_transfers() respectively)

    Without doing so, it can lead to both the interrupt handler and the
    pullup disable routine both writing to the GEVNTCOUNT register, which
    will cause an incorrect count being read from future interrupts.

    Fixes: ae7e86108b12 ("usb: dwc3: Stop active transfers before halting the controller")
    Signed-off-by: Wesley Cheng <wcheng@codeaurora.org>
    Link: https://lore.kernel.org/r/1621571037-1424-1-git-send-email-wcheng@codeaurora.org
```

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

commit ed57e68f8636801fa01d09659841ba982b0c50bf
Author: Sanket Parmar <sparmar@cadence.com>
Date:   Mon May 17 17:05:12 2021 +0200

    usb: cdns3: Enable TDL_CHK only for OUT ep

    commit d6eef886903c4bb5af41b9a31d4ba11dc7a6f8e8 upstream.

    ZLP gets stuck if TDL_CHK bit is set and TDL_FROM_TRB is used
    as TDL source for IN endpoints. To fix it, TDL_CHK is only
    enabled for OUT endpoints.

    Fixes: 7733f6c32e36 ("usb: cdns3: Add Cadence USB3 DRD Driver")
    Reported-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Sanket Parmar <sparmar@cadence.com>
    Link: https://lore.kernel.org/r/1621263912-13175-1-git-send-email-sparmar@cadence.com
    Signed-off-by: Peter Chen <peter.chen@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9a801478908ceaf486905548b7b6419cf6aa9c41
Author: Chen Li <chenli@uniontech.com>
Date:   Tue Apr 27 15:17:45 2021 +0800

    phy: usb: Fix misuse of IS_ENABLED

    commit 7c2fc79250cafa1a29befeb60163028ec4720814 upstream.

    While IS_ENABLED() is perfectly fine for CONFIG_* symbols, it is not
    for other symbols such as __BIG_ENDIAN that is provided directly by
    the compiler.

    Switch to use CONFIG_CPU_BIG_ENDIAN instead of __BIG_ENDIAN.

    Signed-off-by: Chen Li <chenli@uniontech.com>
    Reviewed-by: Al Cooper <alcooperx@gmail.com>
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Fixes: 94583a41047e ("phy: usb: Restructure in preparation for adding 7216 USB support")
    Link: https://lore.kernel.org/r/87czuggpra.wl-chenli@uniontech.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b7611c3679fe3fec8321a1fb4cb44abad51123e3
Author: Axel Lin <axel.lin@ingics.com>
Date:   Sun May 30 20:41:00 2021 +0800

    regulator: rtmv20: Fix .set_current_limit/.get_current_limit callbacks

    commit 86ab21cc39e6b99b7065ab9008c90bec5dec535a upstream.

    Current code does not set .curr_table and .n_linear_ranges settings,
    so it cannot use the regulator_get/set_current_limit_regmap helpers.
    If we setup the curr_table, it will has 200 entries.
    Implement customized .set_current_limit/.get_current_limit callbacks
    instead.

    Fixes: b8c054a5eaf0 ("regulator: rtmv20: Adds support for Richtek RTMV20 load switch
regulator")
    Signed-off-by: Axel Lin <axel.lin@ingics.com>
    Reviewed-by: ChiYuan Huang <cy_huang@richtek.com>
    Link: https://lore.kernel.org/r/20210530124101.477727-1-axel.lin@ingics.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 131503cf85365f9ffd02f6f0b4e75a159b308928
Author: Axel Lin <axel.lin@ingics.com>
Date:   Sat May 29 07:01:47 2021 +0800

    regulator: atc260x: Fix n_voltages and min_sel for pickable linear ranges

    commit 1963fa67d78674a110bc9b2a8b1e226967692f05 upstream.

    The .n_voltages was missed for pickable linear ranges, fix it.
    The min_sel for each pickable range should be starting from 0.
    Also fix atc260x_ldo_voltage_range_sel setting (bit 5 - LDO<N>_VOL_SEL
    in datasheet).

    Fixes: 3b15ccac161a ("regulator: Add regulator driver for ATC260x PMICs")
    Signed-off-by: Axel Lin <axel.lin@ingics.com>
    Reviewed-by: Cristian Ciocaltea <cristian.ciocaltea@gmail.com>
    Link: https://lore.kernel.org/r/20210528230147.363974-1-axel.lin@ingics.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4c445591c1a9d49ba628fae944502607b71a432b
Author: Axel Lin <axel.lin@ingics.com>
Date:   Sun May 23 15:10:45 2021 +0800

    regulator: bd71828: Fix .n_voltages settings

    commit 4c668630bf8ea90a041fc69c9984486e0f56682d upstream.

    Current .n_voltages settings do not cover the latest 2 valid selectors,
    so it fails to set voltage for the hightest voltage support.
    The latest linear range has step_uV = 0, so it does not matter if we
    count the .n_voltages to maximum selector + 1 or the first selector of
    latest linear range + 1.
    To simplify calculating the n_voltages, let's just set the
    .n_voltages to maximum selector + 1.

    Fixes: 522498f8cb8c ("regulator: bd71828: Basic support for ROHM bd71828 PMIC
regulators")
    Signed-off-by: Axel Lin <axel.lin@ingics.com>
    Reviewed-by: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
    Link: https://lore.kernel.org/r/20210523071045.2168904-2-axel.lin@ingics.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f61978e7dda5066207b7098feec0ebbc25e6bc8
Author: Axel Lin <axel.lin@ingics.com>
Date:   Fri May 21 15:30:20 2021 +0800

    regulator: scmi: Fix off-by-one for linear regulators .n_voltages setting

    commit 36cb555fae0875d5416e8514a84a427bec6e4cda upstream.

    For linear regulators, the .n_voltages is (max_uv - min_uv) / uv_step + 1.

    Fixes: 0fbeae70ee7c ("regulator: add SCMI driver")
    Signed-off-by: Axel Lin <axel.lin@ingics.com>
    Reviewed-by: Cristian Marussi <cristian.marussi@arm.com>
    Link: https://lore.kernel.org/r/20210521073020.1944981-1-axel.lin@ingics.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 30db23bc246bf9a436742cf45b8aa1c8d6f2bacc
Author: Axel Lin <axel.lin@ingics.com>
Date:   Thu May 20 19:18:11 2021 +0800
```

```
      regulator: fixed: Ensure enable_counter is correct if reg_domain_disable fails

      commit 855bfff9d623e7aff6556bfb6831d324dec8d96a upstream.

      dev_pm_genpd_set_performance_state() may fail, so had better to check it's
      return value before decreasing priv->enable_counter.

      Fixes: bf3a28cf4241 ("regulator: fixed: support using power domain for enable/disable")
      Signed-off-by: Axel Lin <axel.lin@ingics.com>
      Link: https://lore.kernel.org/r/20210520111811.1806293-1-axel.lin@ingics.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit b4b743881f6820cb4ede49d71d4ff4a4a8d920f5
  Author: Axel Lin <axel.lin@ingics.com>
  Date:   Mon May 17 18:53:24 2021 +0800

      regulator: fan53880: Fix missing n_voltages setting

      commit 34991ee96fd8477479dd15adadceb6b28b30d9b0 upstream.

      Fixes: e6dea51e2d41 ("regulator: fan53880: Add initial support")
      Signed-off-by: Axel Lin <axel.lin@ingics.com>
      Acked-by: Christoph Fritz <chf.fritz@googlemail.com>
      Link: https://lore.kernel.org/r/20210517105325.1227393-1-axel.lin@ingics.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit f52a932f48971dda48b4b74d466d9a00a81f152c
  Author: Axel Lin <axel.lin@ingics.com>
  Date:   Mon May 17 13:27:21 2021 +0800

      regulator: da9121: Return REGULATOR_MODE_INVALID for invalid mode

      commit 0b1e552673724832b08d49037cdeeac634a3b319 upstream.

      -EINVAL is not a valid return value for .of_map_mode, return
      REGULATOR_MODE_INVALID instead.

      Fixes: 65ac97042d4e ("regulator: da9121: add mode support")
      Signed-off-by: Axel Lin <axel.lin@ingics.com>
      Acked-by: Adam Ward <Adam.Ward.opensource@diasemi.com>
      Link: https://lore.kernel.org/r/20210517052721.1063375-1-axel.lin@ingics.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 76d8e9b60bca192b10b29587268118e16365b4bc
  Author: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
  Date:   Wed Jun 2 08:45:58 2021 +0300

      regulator: bd718x7: Fix the BUCK7 voltage setting on BD71837

      commit bc537e65b09a05923f98a31920d1ab170e648dba upstream.

      Changing the BD71837 voltages for other regulators except the first 4 BUCKs
      should be forbidden when the regulator is enabled. There may be out-of-spec
      voltage spikes if the voltage of these "non DVS" bucks is changed when
      enabled. This restriction was accidentally removed when the LDO voltage
      change was allowed for BD71847. (It was not noticed that the BD71837
      BUCK7 used same voltage setting function as LDOs).

      Additionally this bug causes incorrect voltage monitoring register access.
      The voltage change function accidentally used for bd71837 BUCK7 is
      intended to only handle LDO voltage changes. A BD71847 LDO specific
      voltage monitoring disabling code gets executed on BD71837 and register
```

```
    offsets are wrongly calculated as regulator is assumed to be an LDO.

    Prevent the BD71837 BUCK7 voltage change when BUCK7 is enabled by using
    the correct voltage setting operation.

    Fixes: 9bcbabafa19b ("regulator: bd718x7: remove voltage change restriction from BD71847
LDOs")
    Signed-off-by: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
    Link:
https://lore.kernel.org/r/bd8c00931421fafa57e3fdf46557a83075b7cc17.1622610103.git.matti.vait
tinen@fi.rohmeurope.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0d36883c4b336a3401d074908b0c46f9bd17c1ea
Author: Dmitry Osipenko <digetx@gmail.com>
Date:   Mon May 24 01:42:42 2021 +0300

    regulator: max77620: Use device_set_of_node_from_dev()

    commit 6f55c5dd1118b3076d11d9cb17f5c5f4bc3a1162 upstream.

    The MAX77620 driver fails to re-probe on deferred probe because driver
    core tries to claim resources that are already claimed by the PINCTRL
    device. Use device_set_of_node_from_dev() helper which marks OF node as
    reused, skipping erroneous execution of pinctrl_bind_pins() for the PMIC
    device on the re-probe.

    Fixes: aea6cb99703e ("regulator: resolve supply after creating regulator")
    Signed-off-by: Dmitry Osipenko <digetx@gmail.com>
    Link: https://lore.kernel.org/r/20210523224243.13219-2-digetx@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b0dd9b3470fdbf4d7b8e714054ba869c4aeede32
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Thu May 20 01:12:23 2021 +0300

    regulator: core: resolve supply for boot-on/always-on regulators

    commit 98e48cd9283dbac0e1445ee780889f10b3d1db6a upstream.

    For the boot-on/always-on regulators the set_machine_constrainst() is
    called before resolving rdev->supply. Thus the code would try to enable
    rdev before enabling supplying regulator. Enforce resolving supply
    regulator before enabling rdev.

    Fixes: aea6cb99703e ("regulator: resolve supply after creating regulator")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Link: https://lore.kernel.org/r/20210519221224.2868496-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b83c739ad80093309dad82e1189b47eaf14d1439
Author: Kyle Tso <kyletso@google.com>
Date:   Sun Jun 6 16:14:52 2021 +0800

    usb: typec: tcpm: Do not finish VDM AMS for retrying Responses

    commit 5ab14ab1f2db24ffae6c5c39a689660486962e6e upstream.

    If the VDM responses couldn't be sent successfully, it doesn't need to
    finish the AMS until the retry count reaches the limit.

    Fixes: 0908c5aca31e ("usb: typec: tcpm: AMS and Collision Avoidance")
```

```
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Cc: stable <stable@vger.kernel.org>
    Acked-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210606081452.764032-1-kyletso@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cd32d4ee41712d6ff9b56b6d4aea74bce15f00fd
Author: Li Jun <jun.li@nxp.com>
Date:   Wed Jun 2 17:57:08 2021 +0800

    usb: typec: tcpm: cancel frs hrtimer when unregister tcpm port

    commit 7ade4805e296c8d1e40c842395bbe478c7210555 upstream.

    Like the state_machine_timer, we should also cancel possible pending
    frs hrtimer when unregister tcpm port.

    Fixes: 8dc4bd073663 ("usb: typec: tcpm: Add support for Sink Fast Role SWAP(FRS)")
    Cc: stable <stable@vger.kernel.org>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Li Jun <jun.li@nxp.com>
    Link: https://lore.kernel.org/r/1622627829-11070-2-git-send-email-jun.li@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d0a06696a8a4d99f649240b6f9b8a2e55452ecf5
Author: Li Jun <jun.li@nxp.com>
Date:   Wed Jun 2 17:57:07 2021 +0800

    usb: typec: tcpm: cancel vdm and state machine hrtimer when unregister tcpm port

    commit 3a13ff7ef4349d70d1d18378d661117dd5af8efe upstream.

    A pending hrtimer may expire after the kthread_worker of tcpm port
    is destroyed, see below kernel dump when do module unload, fix it
    by cancel the 2 hrtimers.

    [  111.517018] Unable to handle kernel paging request at virtual address
ffff8000118cb880
    [  111.518786] blk_update_request: I/O error, dev sda, sector 60061185 op 0x0:(READ)
flags 0x0 phys_seg 1 prio class 0
    [  111.526594] Mem abort info:
    [  111.526597]   ESR = 0x96000047
    [  111.526600]   EC = 0x25: DABT (current EL), IL = 32 bits
    [  111.526604]   SET = 0, FnV = 0
    [  111.526607]   EA = 0, S1PTW = 0
    [  111.526610] Data abort info:
    [  111.526612]   ISV = 0, ISS = 0x00000047
    [  111.526615]   CM = 0, WnR = 1
    [  111.526619] swapper pgtable: 4k pages, 48-bit VAs, pgdp=0000000041d75000
    [  111.526623] [ffff8000118cb880] pgd=10000001bffff003, p4d=10000001bffff003,
pud=10000001bfffe003, pmd=10000001bfffa003, pte=0000000000000000
    [  111.526642] Internal error: Oops: 96000047 [#1] PREEMPT SMP
    [  111.526647] Modules linked in: dwc3_imx8mp dwc3 phy_fsl_imx8mq_usb [last unloaded:
tcpci]
    [  111.526663] CPU: 0 PID: 0 Comm: swapper/0 Not tainted 5.13.0-rc4-00927-gebbe9dbd802c-
dirty #36
    [  111.526670] Hardware name: NXP i.MX8MPlus EVK board (DT)
    [  111.526674] pstate: 800000c5 (Nzcv daIF -PAN -UAO -TCO BTYPE=--)
    [  111.526681] pc : queued_spin_lock_slowpath+0x1a0/0x390
    [  111.526695] lr : _raw_spin_lock_irqsave+0x88/0xb4
    [  111.526703] sp : ffff800010003e20
    [  111.526706] x29: ffff800010003e20 x28: ffff00017f380180
    [  111.537156] buffer_io_error: 6 callbacks suppressed
    [  111.537162] Buffer I/O error on dev sda1, logical block 60040704, async page read
```

```
    [  111.539932]  x27: ffff00017f3801c0
    [  111.539938] x26: ffff800010ba2490 x25: 0000000000000000 x24: 0000000000000001
    [  111.543025] blk_update_request: I/O error, dev sda, sector 60061186 op 0x0:(READ)
flags 0x0 phys_seg 7 prio class 0
    [  111.548304]
    [  111.548306] x23: 00000000000000c0 x22: ffff0000c2a9f184 x21: ffff00017f380180
    [  111.551374] Buffer I/O error on dev sda1, logical block 60040705, async page read
    [  111.554499]
    [  111.554503] x20: ffff0000c5f14210 x19: 00000000000000c0 x18: 0000000000000000
    [  111.557391] Buffer I/O error on dev sda1, logical block 60040706, async page read
    [  111.561218]
    [  111.561222] x17: 0000000000000000 x16: 0000000000000000 x15: 0000000000000000
    [  111.564205] Buffer I/O error on dev sda1, logical block 60040707, async page read
    [  111.570887] x14: 00000000000000f5 x13: 0000000000000001 x12: 0000000000000040
    [  111.570902] x11: ffff0000c05ac6d8
    [  111.583420] Buffer I/O error on dev sda1, logical block 60040708, async page read
    [  111.588978]  x10: 0000000000000000 x9 : 0000000000040000
    [  111.588988] x8 : 0000000000000000
    [  111.597173] Buffer I/O error on dev sda1, logical block 60040709, async page read
    [  111.605766]  x7 : ffff00017f384880 x6 : ffff8000118cb880
    [  111.605777] x5 : ffff00017f384880
    [  111.611094] Buffer I/O error on dev sda1, logical block 60040710, async page read
    [  111.617086]  x4 : 0000000000000000 x3 : ffff0000c2a9f184
    [  111.617096] x2 : ffff8000118cb880
    [  111.622242] Buffer I/O error on dev sda1, logical block 60040711, async page read
    [  111.626927]  x1 : ffff8000118cb880 x0 : ffff00017f384888
    [  111.626938] Call trace:
    [  111.626942]  queued_spin_lock_slowpath+0x1a0/0x390
    [  111.795809]  kthread_queue_work+0x30/0xc0
    [  111.799828]  state_machine_timer_handler+0x20/0x30
    [  111.804624]  __hrtimer_run_queues+0x140/0x1e0
    [  111.808990]  hrtimer_interrupt+0xec/0x2c0
    [  111.813004]  arch_timer_handler_phys+0x38/0x50
    [  111.817456]  handle_percpu_devid_irq+0x88/0x150
    [  111.821991]  __handle_domain_irq+0x80/0xe0
    [  111.826093]  gic_handle_irq+0xc0/0x140
    [  111.829848]  el1_irq+0xbc/0x154
    [  111.832991]  arch_cpu_idle+0x1c/0x2c
    [  111.836572]  default_idle_call+0x24/0x6c
    [  111.840497]  do_idle+0x238/0x2ac
    [  111.843729]  cpu_startup_entry+0x2c/0x70
    [  111.847657]  rest_init+0xdc/0xec
    [  111.850890]  arch_call_rest_init+0x14/0x20
    [  111.854988]  start_kernel+0x508/0x540
    [  111.858659] Code: 910020e0 8b0200c2 f861d884 aa0203e1 (f8246827)
    [  111.864760] ---[ end trace 308b9a4a3dcb73ac ]---
    [  111.869381] Kernel panic - not syncing: Oops: Fatal exception in interrupt
    [  111.876258] SMP: stopping secondary CPUs
    [  111.880185] Kernel Offset: disabled
    [  111.883673] CPU features: 0x00001001,20000846
    [  111.888031] Memory Limit: none
    [  111.891090] ---[ end Kernel panic - not syncing: Oops: Fatal exception in interrupt
]---

    Fixes: 3ed8e1c2ac99 ("usb: typec: tcpm: Migrate workqueue to RT priority for processing
events")
    Cc: stable <stable@vger.kernel.org>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Li Jun <jun.li@nxp.com>
    Link: https://lore.kernel.org/r/1622627829-11070-1-git-send-email-jun.li@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b2d5b310e115bc928ead124a05a620dff9dce8eb
Author: Kyle Tso <kyletso@google.com>
Date:   Tue Jun 1 00:49:28 2021 +0800
```

```
    usb: typec: tcpm: Properly handle Alert and Status Messages

    commit 063933f47a7af01650af9c4fbcc5831f1c4eb7d9 upstream.

    When receiving Alert Message, if it is not unexpected but is
    unsupported for some reason, the port should return Not_Supported
    Message response.

    Also, according to PD3.0 Spec 6.5.2.1.4 Event Flags Field, the
    OTP/OVP/OCP flags in the Event Flags field in Status Message no longer
    require Get_PPS_Status Message to clear them. Thus remove it when
    receiving Status Message with those flags being set.

    In addition, add the missing AMS operations for Status Message.

    Fixes: 64f7c494a3c0 ("typec: tcpm: Add support for sink PPS related messages")
    Fixes: 0908c5aca31e ("usb: typec: tcpm: AMS and Collision Avoidance")
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210531164928.2368606-1-kyletso@google.com
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ca6bc277430d90375452b60b047763a090b7673e
Author: Maciej Żenczykowski <maze@google.com>
Date:   Tue Jun 8 19:44:59 2021 -0700

    usb: fix various gadget panics on 10gbps cabling

    commit 032e288097a553db5653af552dd8035cd2a0ba96 upstream.

    usb_assign_descriptors() is called with 5 parameters,
    the last 4 of which are the usb_descriptor_header for:
      full-speed (USB1.1 - 12Mbps [including USB1.0 low-speed @ 1.5Mbps),
      high-speed (USB2.0 - 480Mbps),
      super-speed (USB3.0 - 5Gbps),
      super-speed-plus (USB3.1 - 10Gbps).

    The differences between full/high/super-speed descriptors are usually
    substantial (due to changes in the maximum usb block size from 64 to 512
    to 1024 bytes and other differences in the specs), while the difference
    between 5 and 10Gbps descriptors may be as little as nothing
    (in many cases the same tuning is simply good enough).

    However if a gadget driver calls usb_assign_descriptors() with
    a NULL descriptor for super-speed-plus and is then used on a max 10gbps
    configuration, the kernel will crash with a null pointer dereference,
    when a 10gbps capable device port + cable + host port combination shows up.
    (This wouldn't happen if the gadget max-speed was set to 5gbps, but
    it of course defaults to the maximum, and there's no real reason to
    artificially limit it)

    The fix is to simply use the 5gbps descriptor as the 10gbps descriptor,
    if a 10gbps descriptor wasn't provided.

    Obviously this won't fix the problem if the 5gbps descriptor is also
    NULL, but such cases can't be so trivially solved (and any such gadgets
    are unlikely to be used with USB3 ports any way).

    Cc: Felipe Balbi <balbi@kernel.org>
    Cc: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Maciej Żenczykowski <maze@google.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210609024459.1126080-1-zenczykowski@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit f17aae7c4009160f0630a91842a281773976a5bc
Author: Maciej Żenczykowski <maze@google.com>
Date:   Mon Jun 7 21:41:41 2021 -0700

    usb: fix various gadgets null ptr deref on 10gbps cabling.

    commit 90c4d05780d47e14a50e11a7f17373104cd47d25 upstream.

    This avoids a null pointer dereference in
    f_{ecm,eem,hid,loopback,printer,rndis,serial,sourcesink,subset,tcm}
    by simply reusing the 5gbps config for 10gbps.

    Fixes: eaef50c76057 ("usb: gadget: Update usb_assign_descriptors for SuperSpeedPlus")
    Cc: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Cc: Felipe Balbi <balbi@kernel.org>
    Cc: Gustavo A. R. Silva <gustavoars@kernel.org>
    Cc: Lorenzo Colitti <lorenzo@google.com>
    Cc: Martin K. Petersen <martin.petersen@oracle.com>
    Cc: Michael R Sweet <msweet@msweet.org>
    Cc: Mike Christie <michael.christie@oracle.com>
    Cc: Pawel Laszczak <pawell@cadence.com>
    Cc: Peter Chen <peter.chen@nxp.com>
    Cc: Sudhakar Panneerselvam <sudhakar.panneerselvam@oracle.com>
    Cc: Wei Ming Chen <jj251510319013@gmail.com>
    Cc: Will McVicker <willmcvicker@google.com>
    Cc: Zqiang <qiang.zhang@windriver.com>
    Reviewed-By: Lorenzo Colitti <lorenzo@google.com>
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Maciej Żenczykowski <maze@google.com>
    Link: https://lore.kernel.org/r/20210608044141.3898496-1-zenczykowski@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0ee582bcc78615a95b45a793815909c30ca0d981
Author: Linyu Yuan <linyyuan@codeaurora.com>
Date:   Wed Jun 9 07:35:47 2021 +0800

    usb: gadget: eem: fix wrong eem header operation

    commit 305f670846a31a261462577dd0b967c4fa796871 upstream.

    when skb_clone() or skb_copy_expand() fail,
    it should pull skb with lengh indicated by header,
    or not it will read network data and check it as header.

    Cc: <stable@vger.kernel.org>
    Signed-off-by: Linyu Yuan <linyyuan@codeaurora.com>
    Link: https://lore.kernel.org/r/20210608233547.3767-1-linyyuan@codeaurora.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ae746b6f4ce619cf4032fd798a232b010907a397
Author: Pawel Laszczak <pawell@cadence.com>
Date:   Wed May 26 08:05:27 2021 +0200

    usb: cdnsp: Fix deadlock issue in cdnsp_thread_irq_handler

    commit a9aecef198faae3240921b707bc09b602e966fce upstream.

    Patch fixes the following critical issue caused by deadlock which has been
    detected during testing NCM class:

    smp: csd: Detected non-responsive CSD lock (#1) on CPU#0
    smp:     csd: CSD lock (#1) unresponsive.
    ....
    RIP: 0010:native_queued_spin_lock_slowpath+0x61/0x1d0
```

```
      RSP: 0018:ffffbc494011cde0 EFLAGS: 00000002
      RAX: 0000000000000101 RBX: ffff9ee8116b4a68 RCX: 0000000000000000
      RDX: 0000000000000000 RSI: 0000000000000000 RDI: ffff9ee8116b4658
      RBP: ffffbc494011cde0 R08: 0000000000000001 R09: 0000000000000000
      R10: ffff9ee8116b4670 R11: 0000000000000000 R12: ffff9ee8116b4658
      R13: ffff9ee8116b4670 R14: 0000000000000246 R15: ffff9ee8116b4658
      CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
      CR2: 00007f7bcc41a830 CR3: 000000007a612003 CR4: 00000000001706e0
      Call Trace:
       <IRQ>
       do_raw_spin_lock+0xc0/0xd0
       _raw_spin_lock_irqsave+0x95/0xa0
       cdnsp_gadget_ep_queue.cold+0x88/0x107 [cdnsp_udc_pci]
       usb_ep_queue+0x35/0x110
       eth_start_xmit+0x220/0x3d0 [u_ether]
       ncm_tx_timeout+0x34/0x40 [usb_f_ncm]
       ? ncm_free_inst+0x50/0x50 [usb_f_ncm]
       __hrtimer_run_queues+0xac/0x440
       hrtimer_run_softirq+0x8c/0xb0
       __do_softirq+0xcf/0x428
       asm_call_irq_on_stack+0x12/0x20
       </IRQ>
       do_softirq_own_stack+0x61/0x70
       irq_exit_rcu+0xc1/0xd0
       sysvec_apic_timer_interrupt+0x52/0xb0
       asm_sysvec_apic_timer_interrupt+0x12/0x20
      RIP: 0010:do_raw_spin_trylock+0x18/0x40
      RSP: 0018:ffffbc494138bda8 EFLAGS: 00000246
      RAX: 0000000000000000 RBX: ffff9ee8116b4658 RCX: 0000000000000000
      RDX: 0000000000000001 RSI: 0000000000000000 RDI: ffff9ee8116b4658
      RBP: ffffbc494138bda8 R08: 0000000000000001 R09: 0000000000000000
      R10: ffff9ee8116b4670 R11: 0000000000000000 R12: ffff9ee8116b4658
      R13: ffff9ee8116b4670 R14: ffff9ee7b5c73d80 R15: ffff9ee8116b4000
       _raw_spin_lock+0x3d/0x70
       ? cdnsp_thread_irq_handler.cold+0x32/0x112c [cdnsp_udc_pci]
       cdnsp_thread_irq_handler.cold+0x32/0x112c [cdnsp_udc_pci]
       ? cdnsp_remove_request+0x1f0/0x1f0 [cdnsp_udc_pci]
       ? cdnsp_thread_irq_handler+0x5/0xa0 [cdnsp_udc_pci]
       ? irq_thread+0xa0/0x1c0
       irq_thread_fn+0x28/0x60
       irq_thread+0x105/0x1c0
       ? __kthread_parkme+0x42/0x90
       ? irq_forced_thread_fn+0x90/0x90
       ? wake_threads_waitq+0x30/0x30
       ? irq_thread_check_affinity+0xe0/0xe0
       kthread+0x12a/0x160
       ? kthread_park+0x90/0x90
       ret_from_fork+0x22/0x30

      The root cause of issue is spin_lock/spin_unlock instruction instead
      spin_lock_irqsave/spin_lock_irqrestore in cdnsp_thread_irq_handler
      function.

      Cc: stable@vger.kernel.org
      Fixes: 3d82904559f4 ("usb: cdnsp: cdns3 Add main part of Cadence USBSSP DRD Driver")
      Signed-off-by: Pawel Laszczak <pawell@cadence.com>
      Link: https://lore.kernel.org/r/20210526060527.7197-1-pawell@gli-login.cadence.com
      Signed-off-by: Peter Chen <peter.chen@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 04e24bd0f02b67303936868fa94eddb8f6ad7319
  Author: Johan Hovold <johan@kernel.org>
  Date:   Wed Jun 9 18:15:09 2021 +0200

      USB: serial: cp210x: fix CP2102N-A01 modem control
```

```
    commit 63a8eef70ccb5199534dec56fed9759d214bfe55 upstream.

    CP2102N revision A01 (firmware version <= 1.0.4) has a buggy
    flow-control implementation that uses the ulXonLimit instead of
    ulFlowReplace field of the flow-control settings structure (erratum
    CP2102N_E104).

    A recent change that set the input software flow-control limits
    incidentally broke RTS control for these devices when CRTSCTS is not set
    as the new limits would always enable hardware flow control.

    Fix this by explicitly disabling flow control for the buggy firmware
    versions and only updating the input software flow-control limits when
    IXOFF is requested. This makes sure that the terminal settings matches
    the default zero ulXonLimit (ulFlowReplace) for these devices.

    Link: https://lore.kernel.org/r/20210609161509.9459-1-johan@kernel.org
    Reported-by: David Frey <dpfrey@gmail.com>
    Reported-by: Alex Villacís Lasso <a_villacis@palosanto.com>
    Tested-by: Alex Villacís Lasso <a_villacis@palosanto.com>
    Fixes: f61309d9c96a ("USB: serial: cp210x: set IXOFF thresholds")
    Cc: stable@vger.kernel.org      # 5.12
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 802fcbd99ce5a793016ca17ca16fa5c3583149e2
Author: Stefan Agner <stefan@agner.ch>
Date:   Fri May 28 22:39:31 2021 +0200

    USB: serial: cp210x: fix alternate function for CP2102N QFN20

    commit 6f7ec77cc8b64ff5037c1945e4650c65c458037d upstream.

    The QFN20 part has a different GPIO/port function assignment. The
    configuration struct bit field ordered as TX/RX/RS485/WAKEUP/CLK
    which exactly matches GPIO0-3 for QFN24/28. However, QFN20 has a
    different GPIO to primary function assignment.

    Special case QFN20 to follow to properly detect which GPIOs are
    available.

    Signed-off-by: Stefan Agner <stefan@agner.ch>
    Link:
https://lore.kernel.org/r/51830b2b24118eb0f77c5c9ac64ffb2f519dbb1d.1622218300.git.stefan@agn
er.ch
    Fixes: c8acfe0aadbe ("USB: serial: cp210x: implement GPIO support for CP2102N")
    Cc: stable@vger.kernel.org      # 4.19
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0337fee5792a2d0acb6707690c971eaaf634f1c8
Author: Johan Hovold <johan@kernel.org>
Date:   Mon May 24 11:17:05 2021 +0200

    USB: serial: quatech2: fix control-request directions

    commit eb8dbe80326c3d44c1e38ee4f40e0d8d3e06f2d0 upstream.

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

    Fix the three requests which erroneously used usb_rcvctrlpipe().
```

```
    Fixes: f7a33e608d9a ("USB: serial: add quatech2 usb to serial driver")
    Cc: stable@vger.kernel.org      # 3.5
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit db5cd62db582a0c94f4637194b7b3511e2f06b34
Author: Alexandre GRIVEAUX <agriveaux@deutnet.info>
Date:   Sun May 23 18:35:21 2021 +0200

    USB: serial: omninet: add device id for Zyxel Omni 56K Plus

    commit fc0b3dc9a11771c3919eaaaf9d649138b095aa0f upstream.

    Add device id for Zyxel Omni 56K Plus modem, this modem include:

    USB chip:
    NetChip
    NET2888

    Main chip:
    901041A
    F721501APGF

    Another modem using the same chips is the Zyxel Omni 56K DUO/NEO,
    could be added with the right USB ID.

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

commit ae923b092c6906c244102e7f5c9347b7a56cabc5
Author: George McCollister <george.mccollister@gmail.com>
Date:   Thu Jun 3 19:32:08 2021 -0500

    USB: serial: ftdi_sio: add NovaTech OrionMX product ID

    commit bc96c72df33ee81b24d87eab953c73f7bcc04f29 upstream.

    Add PID for the NovaTech OrionMX so it can be automatically detected.

    Signed-off-by: George McCollister <george.mccollister@gmail.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Johan Hovold <johan@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3d472673eb2c5b706ff6e70ce7aa1d01f5e633cc
Author: Wesley Cheng <wcheng@codeaurora.org>
Date:   Fri May 21 17:44:21 2021 -0700

    usb: gadget: f_fs: Ensure io_completion_wq is idle during unbind

    commit 6fc1db5e6211e30fbb1cee8d7925d79d4ed2ae14 upstream.

    During unbind, ffs_func_eps_disable() will be executed, resulting in
    completion callbacks for any pending USB requests.  When using AIO,
    irrespective of the completion status, io_data work is queued to
    io_completion_wq to evaluate and handle the completed requests.  Since
    work runs asynchronously to the unbind() routine, there can be a
    scenario where the work runs after the USB gadget has been fully
    removed, resulting in accessing of a resource which has been already
    freed. (i.e. usb_ep_free_request() accessing the USB ep structure)

    Explicitly drain the io_completion_wq, instead of relying on the
    destroy_workqueue() (in ffs_data_put()) to make sure no pending
```

```
    completion work items are running.

    Signed-off-by: Wesley Cheng <wcheng@codeaurora.org>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/1621644261-1236-1-git-send-email-wcheng@codeaurora.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cad7ce2b91b349b6d40f492d32f754f7398dab88
Author: Andy Shevchenko <andy.shevchenko@gmail.com>
Date:   Mon Jun 7 23:50:06 2021 +0300

    usb: typec: intel_pmc_mux: Add missed error check for devm_ioremap_resource()

    commit 843fabdd7623271330af07f1b7fbd7fabe33c8de upstream.

    devm_ioremap_resource() can return an error, add missed check for it.

    Fixes: 43d596e32276 ("usb: typec: intel_pmc_mux: Check the port status before connect")
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210607205007.71458-2-andy.shevchenko@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 49d9e9ef4c75fe5ec1022089b3f0cf501317a6d7
Author: Andy Shevchenko <andy.shevchenko@gmail.com>
Date:   Mon Jun 7 23:50:05 2021 +0300

    usb: typec: intel_pmc_mux: Put fwnode in error case during ->probe()

    commit 1a85b350a7741776a406005b943e3dec02c424ed upstream.

    device_get_next_child_node() bumps a reference counting of a returned variable.
    We have to balance it whenever we return to the caller.

    Fixes: 6701adfa9693 ("usb: typec: driver for Intel PMC mux control")
    Cc: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210607205007.71458-1-andy.shevchenko@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 02cd2a5cd4bc1f38adf558b35cf3643781c7044d
Author: Mayank Rana <mrana@codeaurora.org>
Date:   Wed Jun 9 00:35:35 2021 -0700

    usb: typec: ucsi: Clear PPM capability data in ucsi_init() error path

    commit f247f0a82a4f8c3bfed178d8fd9e069d1424ee4e upstream.

    If ucsi_init() fails for some reason (e.g. ucsi_register_port()
    fails or general communication failure to the PPM), particularly at
    any point after the GET_CAPABILITY command had been issued, this
    results in unwinding the initialization and returning an error.
    However the ucsi structure's ucsi_capability member retains its
    current value, including likely a non-zero num_connectors.
    And because ucsi_init() itself is done in a workqueue a UCSI
    interface driver will be unaware that it failed and may think the
    ucsi_register() call was completely successful.  Later, if
    ucsi_unregister() is called, due to this stale ucsi->cap value it
    would try to access the items in the ucsi->connector array which
    might not be in a proper state or not even allocated at all and
    results in NULL or invalid pointer dereference.
```

```
    Fix this by clearing the ucsi->cap value to 0 during the error
    path of ucsi_init() in order to prevent a later ucsi_unregister()
    from entering the connector cleanup loop.

    Fixes: c1b0bc2dabfa ("usb: typec: Add support for UCSI interface")
    Cc: stable@vger.kernel.org
    Acked-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Signed-off-by: Mayank Rana <mrana@codeaurora.org>
    Signed-off-by: Jack Pham <jackp@codeaurora.org>
    Link: https://lore.kernel.org/r/20210609073535.5094-1-jackp@codeaurora.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 733101e5f0ee4038a5dc709ba82febc241f884be
Author: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
Date:   Wed Jun 9 20:22:02 2021 +0300

    usb: typec: wcove: Use LE to CPU conversion when accessing msg->header

    commit d5ab95da2a41567440097c277c5771ad13928dad upstream.

    As LKP noticed the Sparse is not happy about strict type handling:
       .../typec/tcpm/wcove.c:380:50: sparse:     expected unsigned short [usertype] header
       .../typec/tcpm/wcove.c:380:50: sparse:     got restricted __le16 const [usertype]
header

    Fix this by switching to use pd_header_cnt_le() instead of pd_header_cnt()
    in the affected code.

    Fixes: ae8a2ca8a221 ("usb: typec: Group all TCPCI/TCPM code together")
    Fixes: 3c4fb9f16921 ("usb: typec: wcove: start using tcpm for USB PD support")
    Reported-by: kernel test robot <lkp@intel.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20210609172202.83377-1-andriy.shevchenko@linux.intel.com
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b9f8edb1e44bbc11fcf087a24bf027a54650fb57
Author: Thomas Petazzoni <thomas.petazzoni@bootlin.com>
Date:   Fri May 28 16:04:46 2021 +0200

    usb: musb: fix MUSB_QUIRK_B_DISCONNECT_99 handling

    commit b65ba0c362be665192381cc59e3ac3ef6f0dd1e1 upstream.

    In commit 92af4fc6ec33 ("usb: musb: Fix suspend with devices
    connected for a64"), the logic to support the
    MUSB_QUIRK_B_DISCONNECT_99 quirk was modified to only conditionally
    schedule the musb->irq_work delayed work.

    This commit badly breaks ECM Gadget on AM335X. Indeed, with this
    commit, one can observe massive packet loss:

    $ ping 192.168.0.100
    ...
    15 packets transmitted, 3 received, 80% packet loss, time 14316ms

    Reverting this commit brings back a properly functioning ECM
    Gadget. An analysis of the commit seems to indicate that a mistake was
    made: the previous code was not falling through into the
    MUSB_QUIRK_B_INVALID_VBUS_91, but now it is, unless the condition is
    taken.

    Changing the logic to be as it was before the problematic commit *and*
```

```
    only conditionally scheduling musb->irq_work resolves the regression:

    $ ping 192.168.0.100
    ...
    64 packets transmitted, 64 received, 0% packet loss, time 64475ms

    Fixes: 92af4fc6ec33 ("usb: musb: Fix suspend with devices connected for a64")
    Cc: stable@vger.kernel.org
    Tested-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Tested-by: Drew Fustini <drew@beagleboard.org>
    Acked-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Thomas Petazzoni <thomas.petazzoni@bootlin.com>
    Link: https://lore.kernel.org/r/20210528140446.278076-1-thomas.petazzoni@bootlin.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cb65ff03f087224651718cf481a9b519b6ff08c3
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Thu May 27 10:45:34 2021 -0500

    usb: pci-quirks: disable D3cold on xhci suspend for s2idle on AMD Renoir

    commit d1658268e43980c071dbffc3d894f6f6c4b6732a upstream.

    The XHCI controller is required to enter D3hot rather than D3cold for AMD
    s2idle on this hardware generation.

    Otherwise, the 'Controller Not Ready' (CNR) bit is not being cleared by
    host in resume and eventually this results in xhci resume failures during
    the s2idle wakeup.

    Link: https://lore.kernel.org/linux-usb/1612527609-7053-1-git-send-email-
Prike.Liang@amd.com/
    Suggested-by: Prike Liang <Prike.Liang@amd.com>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Cc: stable <stable@vger.kernel.org> # 5.11+
    Link: https://lore.kernel.org/r/20210527154534.8900-1-mario.limonciello@amd.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 788755756dd4a6aba1de479fec20b0fa600e7f19
Author: Marian-Cristian Rotariu <marian.c.rotariu@gmail.com>
Date:   Tue Jun 8 19:26:50 2021 +0300

    usb: dwc3: ep0: fix NULL pointer exception

    commit d00889080ab60051627dab1d85831cd9db750e2a upstream.

    There is no validation of the index from dwc3_wIndex_to_dep() and we might
    be referring a non-existing ep and trigger a NULL pointer exception. In
    certain configurations we might use fewer eps and the index might wrongly
    indicate a larger ep index than existing.

    By adding this validation from the patch we can actually report a wrong
    index back to the caller.

    In our usecase we are using a composite device on an older kernel, but
    upstream might use this fix also. Unfortunately, I cannot describe the
    hardware for others to reproduce the issue as it is a proprietary
    implementation.

    [   82.958261] Unable to handle kernel NULL pointer dereference at virtual address
00000000000000a4
    [   82.966891] Mem abort info:
    [   82.969663]   ESR = 0x96000006
    [   82.972703]   Exception class = DABT (current EL), IL = 32 bits
    [   82.978603]   SET = 0, FnV = 0
```

```
     [   82.981642]   EA = 0, S1PTW = 0
     [   82.984765] Data abort info:
     [   82.987631]   ISV = 0, ISS = 0x00000006
     [   82.991449]   CM = 0, WnR = 0
     [   82.994409] user pgtable: 4k pages, 39-bit VAs, pgdp = 00000000c6210ccc
     [   83.000999] [0000000000000a4] pgd=0000000053aa5003, pud=0000000053aa5003,
 pmd=0000000000000000
     [   83.009685] Internal error: Oops: 96000006 [#1] PREEMPT SMP
     [   83.026433] Process irq/62-dwc3 (pid: 303, stack limit = 0x000000003985154c)
     [   83.033470] CPU: 0 PID: 303 Comm: irq/62-dwc3 Not tainted 4.19.124 #1
     [   83.044836] pstate: 60000085 (nZCv daIf -PAN -UAO)
     [   83.049628] pc : dwc3_ep0_handle_feature+0x414/0x43c
     [   83.054558] lr : dwc3_ep0_interrupt+0x3b4/0xc94

     ...

     [   83.141788] Call trace:
     [   83.144227]  dwc3_ep0_handle_feature+0x414/0x43c
     [   83.148823]  dwc3_ep0_interrupt+0x3b4/0xc94
     [   83.181546] ---[ end trace aac6b5267d84c32f ]---

     Signed-off-by: Marian-Cristian Rotariu <marian.c.rotariu@gmail.com>
     Cc: stable <stable@vger.kernel.org>
     Link: https://lore.kernel.org/r/20210608162650.58426-1-marian.c.rotariu@gmail.com
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit 4aad390363d2b9b3e92428dd34d27bb7ea8f1ee8
 Author: Jack Pham <jackp@codeaurora.org>
 Date:   Fri May 28 09:04:05 2021 -0700

     usb: dwc3: gadget: Bail from dwc3_gadget_exit() if dwc->gadget is NULL

     commit 03715ea2e3dbbc56947137ce3b4ac18a726b2f87 upstream.

     There exists a possible scenario in which dwc3_gadget_init() can fail:
     during during host -> peripheral mode switch in dwc3_set_mode(), and
     a pending gadget driver fails to bind.  Then, if the DRD undergoes
     another mode switch from peripheral->host the resulting
     dwc3_gadget_exit() will attempt to reference an invalid and dangling
     dwc->gadget pointer as well as call dma_free_coherent() on unmapped
     DMA pointers.

     The exact scenario can be reproduced as follows:
      - Start DWC3 in peripheral mode
      - Configure ConfigFS gadget with FunctionFS instance (or use g_ffs)
      - Run FunctionFS userspace application (open EPs, write descriptors, etc)
      - Bind gadget driver to DWC3's UDC
      - Switch DWC3 to host mode
        => dwc3_gadget_exit() is called. usb_del_gadget() will put the
             ConfigFS driver instance on the gadget_driver_pending_list
      - Stop FunctionFS application (closes the ep files)
      - Switch DWC3 to peripheral mode
        => dwc3_gadget_init() fails as usb_add_gadget() calls
             check_pending_gadget_drivers() and attempts to rebind the UDC
             to the ConfigFS gadget but fails with -19 (-ENODEV) because the
             FFS instance is not in FFS_ACTIVE state (userspace has not
             re-opened and written the descriptors yet, i.e. desc_ready!=0).
      - Switch DWC3 back to host mode
        => dwc3_gadget_exit() is called again, but this time dwc->gadget
             is invalid.

     Although it can be argued that userspace should take responsibility
     for ensuring that the FunctionFS application be ready prior to
     allowing the composite driver bind to the UDC, failure to do so
     should not result in a panic from the kernel driver.
```

```
    Fix this by setting dwc->gadget to NULL in the failure path of
    dwc3_gadget_init() and add a check to dwc3_gadget_exit() to bail out
    unless the gadget pointer is valid.

    Fixes: e81a7018d93a ("usb: dwc3: allocate gadget structure dynamically")
    Cc: <stable@vger.kernel.org>
    Reviewed-by: Peter Chen <peter.chen@kernel.org>
    Signed-off-by: Jack Pham <jackp@codeaurora.org>
    Link: https://lore.kernel.org/r/20210528160405.17550-1-jackp@codeaurora.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2bdf7460af16d5a4b670b6905f882f4f63762c89
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Fri May 21 18:55:50 2021 +0200

    usb: dwc3: meson-g12a: Disable the regulator in the error handling path of the probe

    commit 1d0d3d818eafe1963ec1eaf302175cd14938188e upstream.

    If an error occurs after a successful 'regulator_enable()' call,
    'regulator_disable()' must be called.

    Fix the error handling path of the probe accordingly.

    The remove function doesn't need to be fixed, because the
    'regulator_disable()' call is already hidden in 'dwc3_meson_g12a_suspend()'
    which is called via 'pm_runtime_set_suspended()' in the remove function.

    Fixes: c99993376f72 ("usb: dwc3: Add Amlogic G12A DWC3 glue")
    Reviewed-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Acked-by: Neil Armstrong <narmstrong@baylibre.com>
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Link:
https://lore.kernel.org/r/79df054046224bbb0716a8c5c2082650290eec86.1621616013.git.christophe
.jaillet@wanadoo.fr
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d8dd3754e707104a34f8ec595034d503ea8871a2
Author: Neil Armstrong <narmstrong@baylibre.com>
Date:   Tue Jun 1 10:48:30 2021 +0200

    usb: dwc3-meson-g12a: fix usb2 PHY glue init when phy0 is disabled

    commit 4d2aa178d2ad2fb156711113790dde13e9aa2376 upstream.

    When only PHY1 is used (for example on Odroid-HC4), the regmap init code
    uses the usb2 ports when doesn't initialize the PHY1 regmap entry.

    This fixes:
    Unable to handle kernel NULL pointer dereference at virtual address 0000000000000020
    ...
    pc : regmap_update_bits_base+0x40/0xa0
    lr : dwc3_meson_g12a_usb2_init_phy+0x4c/0xf8
    ...
    Call trace:
    regmap_update_bits_base+0x40/0xa0
    dwc3_meson_g12a_usb2_init_phy+0x4c/0xf8
    dwc3_meson_g12a_usb2_init+0x7c/0xc8
    dwc3_meson_g12a_usb_init+0x28/0x48
    dwc3_meson_g12a_probe+0x298/0x540
    platform_probe+0x70/0xe0
    really_probe+0xf0/0x4d8
    driver_probe_device+0xfc/0x168
```

```
    ...

    Fixes: 013af227f58a97 ("usb: dwc3: meson-g12a: handle the phy and glue registers
separately")
    Reviewed-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Neil Armstrong <narmstrong@baylibre.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210601084830.260196-1-narmstrong@baylibre.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9015140154c5d242ad33885c9e35789859b43c37
Author: Kyle Tso <kyletso@google.com>
Date:   Fri May 28 16:16:13 2021 +0800

    usb: pd: Set PD_T_SINK_WAIT_CAP to 310ms

    commit 6490fa565534fa83593278267785a694fd378a2b upstream.

    Current timer PD_T_SINK_WAIT_CAP is set to 240ms which will violate the
    SinkWaitCapTimer (tTypeCSinkWaitCap 310 - 620 ms) defined in the PD
    Spec if the port is faster enough when running the state machine. Set it
    to the lower bound 310ms to ensure the timeout is in Spec.

    Fixes: f0690a25a140 ("staging: typec: USB Type-C Port Manager (tcpm)")
    Cc: stable <stable@vger.kernel.org>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Kyle Tso <kyletso@google.com>
    Link: https://lore.kernel.org/r/20210528081613.730661-1-kyletso@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8ceb32221c1104409b3311b63b0331b0f2d578d6
Author: Maciej Żenczykowski <maze@google.com>
Date:   Tue Jun 8 01:54:38 2021 -0700

    usb: f_ncm: only first packet of aggregate needs to start timer

    commit 1958ff5ad2d4908b44a72bcf564dfe67c981e7fe upstream.

    The reasoning for this change is that if we already had
    a packet pending, then we also already had a pending timer,
    and as such there is no need to reschedule it.

    This also prevents packets getting delayed 60 ms worst case
    under a tiny packet every 290us transmit load, by keeping the
    timeout always relative to the first queued up packet.
    (300us delay * 16KB max aggregation / 80 byte packet =~ 60 ms)

    As such the first packet is now at most delayed by 300us.

    Under low transmit load, this will simply result in us sending
    a shorter aggregate, as originally intended.

    This patch has the benefit of greatly reducing (by ~10 factor
    with 1500 byte frames aggregated into 16 kiB) the number of
    (potentially pretty costly) updates to the hrtimer.

    Cc: Brooke Basile <brookebasile@gmail.com>
    Cc: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Cc: Felipe Balbi <balbi@kernel.org>
    Cc: Lorenzo Colitti <lorenzo@google.com>
    Signed-off-by: Maciej Żenczykowski <maze@google.com>
    Link: https://lore.kernel.org/r/20210608085438.813960-1-zenczykowski@gmail.com
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 9e5f198de7c8872315cad40317e2d4a9ebfe8472
Author: Maciej Żenczykowski <maze@google.com>
Date:   Mon Jun 7 17:53:44 2021 -0700

    USB: f_ncm: ncm_bitrate (speed) is unsigned

    commit 3370139745853f7826895293e8ac3aec1430508e upstream.

    [  190.544755] configfs-gadget gadget: notify speed -44967296

    This is because 4250000000 - 2**32 is -44967296.

    Fixes: 9f6ce4240a2b ("usb: gadget: f_ncm.c added")
    Cc: Brooke Basile <brookebasile@gmail.com>
    Cc: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Cc: Felipe Balbi <balbi@kernel.org>
    Cc: Lorenzo Colitti <lorenzo@google.com>
    Cc: Yauheni Kaliuta <yauheni.kaliuta@nokia.com>
    Cc: Linux USB Mailing List <linux-usb@vger.kernel.org>
    Acked-By: Lorenzo Colitti <lorenzo@google.com>
    Signed-off-by: Maciej Żenczykowski <maze@google.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210608005344.3762668-1-zenczykowski@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit da7ed080ef4de08213a26a2b725eb02c9b6b28e6
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Jun 4 14:59:43 2021 +0200

    mmc: renesas_sdhi: Fix HS400 on R-Car M3-W+

    commit 6687cd72aa9112a454a4646986e0402dd1b07d0e upstream.

    R-Car M3-W ES3.0 is marketed as R-Car M3-W+ (R8A77961), and has its own
    compatible value "renesas,r8a77961".

    Hence using soc_device_match() with soc_id = "r8a7796" and revision =
    "ES3.*" does not actually match running on an R-Car M3-W+ SoC.

    Fix this by matching with soc_id = "r8a77961" instead.

    Fixes: a38c078fea0b1393 ("mmc: renesas_sdhi: Avoid bad TAP in HS400")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Reviewed-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Link:
https://lore.kernel.org/r/ee8af5d631f5331139ffea714539030d97352e93.1622811525.git.geert+rene
sas@glider.be
    Cc: stable@vger.kernel.org
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 800ad9137bec92ae9d8e64fbe9daea7baa99d6be
Author: Wolfram Sang <wsa+renesas@sang-engineering.com>
Date:   Wed Jun 2 09:34:35 2021 +0200

    mmc: renesas_sdhi: abort tuning when timeout detected

    commit 2c9017d0b5d3fbf17e69577a42d9e610ca122810 upstream.

    We have to bring the eMMC from sending-data state back to transfer state
    once we detected a CRC error (timeout) during tuning. So, send a stop
    command via mmc_abort_tuning().
```

```
    Fixes: 4f11997773b6 ("mmc: tmio: Add tuning support")
    Reported-by Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Signed-off-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Reviewed-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Tested-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Link: https://lore.kernel.org/r/20210602073435.5955-1-wsa+renesas@sang-engineering.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2097283179990102077 8da5901a24276e086c3d
Author: Masami Hiramatsu <mhiramat@kernel.org>
Date:   Thu May 13 12:06:33 2021 +0900

    tools/bootconfig: Fix a build error accroding to undefined fallthrough

    commit 824afd55e95c3cb12c55d297a0ae408be1779cc8 upstream.

    Since the "fallthrough" is defined only in the kernel, building
    lib/bootconfig.c as a part of user-space tools causes a build
    error.

    Add a dummy fallthrough to avoid the build error.

    Link: https://lkml.kernel.org/r/162087519356.442660.11385099982318160180.stgit@devnote2

    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: stable@vger.kernel.org
    Fixes: 4c1ca831adb1 ("Revert "lib: Revert use of fallthrough pseudo-keyword in lib/"")
    Signed-off-by: Masami Hiramatsu <mhiramat@kernel.org>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b16a249eca2230c2cd66fa1d4b94743bd9b6ef92
Author: Liangyan <liangyan.peng@linux.alibaba.com>
Date:   Mon Jun 7 20:57:34 2021 +0800

    tracing: Correct the length check which causes memory corruption

    commit 3e08a9f9760f4a70d633c328a76408e62d6f80a3 upstream.

    We've suffered from severe kernel crashes due to memory corruption on
    our production environment, like,

    Call Trace:
    [1640542.554277] general protection fault: 0000 [#1] SMP PTI
    [1640542.554856] CPU: 17 PID: 26996 Comm: python Kdump: loaded Tainted:G
    [1640542.556629] RIP: 0010:kmem_cache_alloc+0x90/0x190
    [1640542.559074] RSP: 0018:ffffb16faa597df8 EFLAGS: 00010286
    [1640542.559587] RAX: 0000000000000000 RBX: 0000000000400200 RCX:
    0000000006e931bf
    [1640542.560323] RDX: 0000000006e931be RSI: 0000000000400200 RDI:
    ffff9a45ff004300
    [1640542.560996] RBP: 0000000000400200 R08: 0000000000023420 R09:
    0000000000000000
    [1640542.561670] R10: 0000000000000000 R11: 0000000000000000 R12:
    ffffffff9a20608d
    [1640542.562366] R13: ffff9a45ff004300 R14: ffff9a45ff004300 R15:
    696c662f65636976
    [1640542.563128] FS:  00007f45d7c6f740(0000) GS:ffff9a45ff840000(0000)
    knlGS:0000000000000000
    [1640542.563937] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [1640542.564557] CR2: 00007f45d71311a0 CR3: 000000189d63e004 CR4:
    00000000003606e0
```

```
    [1640542.565279] DR0: 0000000000000000 DR1: 0000000000000000 DR2:
    0000000000000000
    [1640542.566069] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7:
    0000000000000400
    [1640542.566742] Call Trace:
    [1640542.567009]  anon_vma_clone+0x5d/0x170
    [1640542.567417]  __split_vma+0x91/0x1a0
    [1640542.567777]  do_munmap+0x2c6/0x320
    [1640542.568128]  vm_munmap+0x54/0x70
    [1640542.569990]  __x64_sys_munmap+0x22/0x30
    [1640542.572005]  do_syscall_64+0x5b/0x1b0
    [1640542.573724]  entry_SYSCALL_64_after_hwframe+0x44/0xa9
    [1640542.575642] RIP: 0033:0x7f45d6e61e27

    James Wang has reproduced it stably on the latest 4.19 LTS.
    After some debugging, we finally proved that it's due to ftrace
    buffer out-of-bound access using a debug tool as follows:
    [   86.775200] BUG: Out-of-bounds write at addr 0xffff88aefe8b7000
    [   86.780806]  no_context+0xdf/0x3c0
    [   86.784327]  __do_page_fault+0x252/0x470
    [   86.788367]  do_page_fault+0x32/0x140
    [   86.792145]  page_fault+0x1e/0x30
    [   86.795576]  strncpy_from_unsafe+0x66/0xb0
    [   86.799789]  fetch_memory_string+0x25/0x40
    [   86.804002]  fetch_deref_string+0x51/0x60
    [   86.808134]  kprobe_trace_func+0x32d/0x3a0
    [   86.812347]  kprobe_dispatcher+0x45/0x50
    [   86.816385]  kprobe_ftrace_handler+0x90/0xf0
    [   86.820779]  ftrace_ops_assist_func+0xa1/0x140
    [   86.825340]  0xffffffffc00750bf
    [   86.828603]  do_sys_open+0x5/0x1f0
    [   86.832124]  do_syscall_64+0x5b/0x1b0
    [   86.835900]  entry_SYSCALL_64_after_hwframe+0x44/0xa9

    commit b220c049d519 ("tracing: Check length before giving out
    the filter buffer") adds length check to protect trace data
    overflow introduced in 0fc1b09ff1ff, seems that this fix can't prevent
    overflow entirely, the length check should also take the sizeof
    entry->array[0] into account, since this array[0] is filled the
    length of trace data and occupy addtional space and risk overflow.

    Link: https://lkml.kernel.org/r/20210607125734.1770447-1-liangyan.peng@linux.alibaba.com

    Cc: stable@vger.kernel.org
    Cc: Ingo Molnar <mingo@redhat.com>
    Cc: Xunlei Pang <xlpang@linux.alibaba.com>
    Cc: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Fixes: b220c049d519 ("tracing: Check length before giving out the filter buffer")
    Reviewed-by: Xunlei Pang <xlpang@linux.alibaba.com>
    Reviewed-by: yinbinbin <yinbinbin@alibabacloud.com>
    Reviewed-by: Wetp Zhang <wetp.zy@linux.alibaba.com>
    Tested-by: James Wang <jnwang@linux.alibaba.com>
    Signed-off-by: Liangyan <liangyan.peng@linux.alibaba.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3e4ddeb68751fb4fb657199aed9cfd5d02796875
Author: Steven Rostedt (VMware) <rostedt@goodmis.org>
Date:   Mon Jun 7 21:39:08 2021 -0400

    ftrace: Do not blindly read the ip address in ftrace_bug()

    commit 6c14133d2d3f768e0a35128faac8aa6ed4815051 upstream.

    It was reported that a bug on arm64 caused a bad ip address to be used for
```

```
    updating into a nop in ftrace_init(), but the error path (rightfully)
    returned -EINVAL and not -EFAULT, as the bug caused more than one error to
    occur. But because -EINVAL was returned, the ftrace_bug() tried to report
    what was at the location of the ip address, and read it directly. This
    caused the machine to panic, as the ip was not pointing to a valid memory
    address.

    Instead, read the ip address with copy_from_kernel_nofault() to safely
    access the memory, and if it faults, report that the address faulted,
    otherwise report what was in that location.

    Link: https://lore.kernel.org/lkml/20210607032329.28671-1-mark-pk.tsai@mediatek.com/

    Cc: stable@vger.kernel.org
    Fixes: 05736a427f7e1 ("ftrace: warn on failure to disable mcount callers")
    Reported-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Tested-by: Mark-PK Tsai <mark-pk.tsai@mediatek.com>
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2145c61cc2b000a522698a9aef3fbaa32ba94027
Author: Alexander Kuznetsov <wwfq@yandex-team.ru>
Date:   Wed Jun 9 10:17:19 2021 +0300

    cgroup1: don't allow '\n' in renaming

    commit b7e24eb1caa5f8da20d405d262dba67943aedc42 upstream.

    cgroup_mkdir() have restriction on newline usage in names:
    $ mkdir $'/sys/fs/cgroup/cpu/test\ntest2'
    mkdir: cannot create directory
    '/sys/fs/cgroup/cpu/test\ntest2': Invalid argument

    But in cgroup1_rename() such check is missed.
    This allows us to make /proc/<pid>/cgroup unparsable:
    $ mkdir /sys/fs/cgroup/cpu/test
    $ mv /sys/fs/cgroup/cpu/test $'/sys/fs/cgroup/cpu/test\ntest2'
    $ echo $$ > $'/sys/fs/cgroup/cpu/test\ntest2'
    $ cat /proc/self/cgroup
    11:pids:/
    10:freezer:/
    9:hugetlb:/
    8:cpuset:/
    7:blkio:/user.slice
    6:memory:/user.slice
    5:net_cls,net_prio:/
    4:perf_event:/
    3:devices:/user.slice
    2:cpu,cpuacct:/test
    test2
    1:name=systemd:/
    0::/

    Signed-off-by: Alexander Kuznetsov <wwfq@yandex-team.ru>
    Reported-by: Andrey Krasichkov <buglloc@yandex-team.ru>
    Acked-by: Dmitry Yakunin <zeil@yandex-team.ru>
    Cc: stable@vger.kernel.org
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ac62404157f14b6ceabef034da12e6c7618bc116
Author: Eric W. Biederman <ebiederm@xmission.com>
Date:   Thu Jun 10 15:11:11 2021 -0500

    coredump: Limit what can interrupt coredumps
```

```
    commit 06af8679449d4ed282df13191fc52d5ba28ec536 upstream.

    Olivier Langlois has been struggling with coredumps being incompletely written in
    processes using io_uring.

    Olivier Langlois <olivier@trillion01.com> writes:
    > io_uring is a big user of task_work and any event that io_uring made a
    > task waiting for that occurs during the core dump generation will
    > generate a TIF_NOTIFY_SIGNAL.
    >
    > Here are the detailed steps of the problem:
    > 1. io_uring calls vfs_poll() to install a task to a file wait queue
    >    with io_async_wake() as the wakeup function cb from io_arm_poll_handler()
    > 2. wakeup function ends up calling task_work_add() with TWA_SIGNAL
    > 3. task_work_add() sets the TIF_NOTIFY_SIGNAL bit by calling
    >    set_notify_signal()

    The coredump code deliberately supports being interrupted by SIGKILL,
    and depends upon prepare_signal to filter out all other signals.   Now
    that signal_pending includes wake ups for TIF_NOTIFY_SIGNAL this hack
    in dump_emitted by the coredump code no longer works.

    Make the coredump code more robust by explicitly testing for all of
    the wakeup conditions the coredump code supports.  This prevents
    new wakeup conditions from breaking the coredump code, as well
    as fixing the current issue.

    The filesystem code that the coredump code uses already limits
    itself to only aborting on fatal_signal_pending.  So it should
    not develop surprising wake-up reasons either.

    v2: Don't remove the now unnecessary code in prepare_signal.

    Cc: stable@vger.kernel.org
    Fixes: 12db8b690010 ("entry: Add support for TIF_NOTIFY_SIGNAL")
    Reported-by: Olivier Langlois <olivier@trillion01.com>
    Signed-off-by: "Eric W. Biederman" <ebiederm@xmission.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7253a49994576d0d61ba5136d6cbd3ac9e819447
Author: Nikolay Borisov <nborisov@suse.com>
Date:   Mon May 31 12:26:01 2021 +0300

    btrfs: promote debugging asserts to full-fledged checks in validate_super

    commit aefd7f7065567a4666f42c0fc8cdb379d2e036bf upstream.

    Syzbot managed to trigger this assert while performing its fuzzing.
    Turns out it's better to have those asserts turned into full-fledged
    checks so that in case buggy btrfs images are mounted the users gets
    an error and mounting is stopped. Alternatively with CONFIG_BTRFS_ASSERT
    disabled such image would have been erroneously allowed to be mounted.

    Reported-by: syzbot+a6bf271c02e4fe66b4e4@syzkaller.appspotmail.com
    CC: stable@vger.kernel.org # 5.4+
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Reviewed-by: Qu Wenruo <wqu@suse.com>
    Signed-off-by: Nikolay Borisov <nborisov@suse.com>
    Reviewed-by: David Sterba <dsterba@suse.com>
    [ add uuids to the messages ]
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit 2818ecb9575313628053b5c481bbeb3a3d8303b6
Author: Ritesh Harjani <riteshh@linux.ibm.com>
Date:   Sun May 30 20:24:05 2021 +0530

    btrfs: return value from btrfs_mark_extent_written() in case of error

    commit e7b2ec3d3d4ebeb4cff7ae45cf430182fa6a49fb upstream.

    We always return 0 even in case of an error in btrfs_mark_extent_written().
    Fix it to return proper error value in case of a failure. All callers
    handle it.

    CC: stable@vger.kernel.org # 4.4+
    Signed-off-by: Ritesh Harjani <riteshh@linux.ibm.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 d7c310163a4b3c715e39e41c2201d5ef140c8947
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Thu May 27 15:27:32 2021 +0900

    btrfs: zoned: fix zone number to sector/physical calculation

    commit 5b434df8778771d181bc19fb4593bca114d1c4eb upstream.

    In btrfs_get_dev_zone_info(), we have "u32 sb_zone" and calculate "sector_t
    sector" by shifting it. But, this "sector" is calculated in 32bit, leading
    it to be 0 for the 2nd superblock copy.

    Since zone number is u32, shifting it to sector (sector_t) or physical
    address (u64) can easily trigger a missing cast bug like this.

    This commit introduces helpers to convert zone number to sector/LBA, so we
    won't fall into the same pitfall again.

    Reported-by: Dmitry Fomichev <Dmitry.Fomichev@wdc.com>
    Fixes: 12659251ca5d ("btrfs: implement log-structured superblock for ZONED mode")
    CC: stable@vger.kernel.org # 5.11+
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.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 fb58673c334e87375cf2ad09a9d12da6e53bd58a
Author: Josef Bacik <josef@toxicpanda.com>
Date:   Wed May 19 17:15:53 2021 -0400

    btrfs: do not write supers if we have an fs error

    commit 165ea85f14831f27fc6fe3b02b35e42e50b9ed94 upstream.

    Error injection testing uncovered a pretty severe problem where we could
    end up committing a super that pointed to the wrong tree roots,
    resulting in transid mismatch errors.

    The way we commit the transaction is we update the super copy with the
    current generations and bytenrs of the important roots, and then copy
    that into our super_for_commit.  Then we allow transactions to continue
    again, we write out the dirty pages for the transaction, and then we
    write the super.  If the write out fails we'll bail and skip writing the
    supers.

    However since we've allowed a new transaction to start, we can have a
```

```
    log attempting to sync at this point, which would be blocked on
    fs_info->tree_log_mutex.  Once the commit fails we're allowed to do the
    log tree commit, which uses super_for_commit, which now points at fs
    tree's that were not written out.

    Fix this by checking BTRFS_FS_STATE_ERROR once we acquire the
    tree_log_mutex.  This way if the transaction commit fails we're sure to
    see this bit set and we can skip writing the super out.  This patch
    fixes this specific transid mismatch error I was seeing with this
    particular error path.

    CC: stable@vger.kernel.org # 5.12+
    Reviewed-by: Filipe Manana <fdmanana@suse.com>
    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: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2c9809adffeeda83b4819110e12b7c15c9f20270
Author: Xiao Ni <xni@redhat.com>
Date:   Fri May 28 14:16:38 2021 +0800

    async_xor: check src_offs is not NULL before updating it

    commit 9be148e408df7d361ec5afd6299b7736ff3928b0 upstream.

    When PAGE_SIZE is greater than 4kB, multiple stripes may share the same
    page. Thus, src_offs is added to async_xor_offs() with array of offsets.
    However, async_xor() passes NULL src_offs to async_xor_offs(). In such
    case, src_offs should not be updated. Add a check before the update.

    Fixes: ceaf2966ab08(async_xor: increase src_offs when dropping destination page)
    Cc: stable@vger.kernel.org # v5.10+
    Reported-by: Oleksandr Shchirskyi <oleksandr.shchirskyi@linux.intel.com>
    Tested-by: Oleksandr Shchirskyi <oleksandr.shchirskyi@intel.com>
    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 555002a840ab88468e252b0eedf0b05e2ce7099c
Author: Coly Li <colyli@suse.de>
Date:   Mon Jun 7 20:50:52 2021 +0800

    bcache: avoid oversized read request in cache missing code path

    commit 41fe8d088e96472f63164e213de44ec77be69478 upstream.

    In the cache missing code path of cached device, if a proper location
    from the internal B+ tree is matched for a cache miss range, function
    cached_dev_cache_miss() will be called in cache_lookup_fn() in the
    following code block,
    [code block 1]
      526         unsigned int sectors = KEY_INODE(k) == s->iop.inode
      527                 ? min_t(uint64_t, INT_MAX,
      528                         KEY_START(k) - bio->bi_iter.bi_sector)
      529                 : INT_MAX;
      530         int ret = s->d->cache_miss(b, s, bio, sectors);

    Here s->d->cache_miss() is the call backfunction pointer initialized as
    cached_dev_cache_miss(), the last parameter 'sectors' is an important
    hint to calculate the size of read request to backing device of the
    missing cache data.

    Current calculation in above code block may generate oversized value of
    'sectors', which consequently may trigger 2 different potential kernel
```

panics by BUG() or BUG_ON() as listed below,

1) BUG_ON() inside bch_btree_insert_key(),
[code block 2]
```
   886         BUG_ON(b->ops->is_extents && !KEY_SIZE(k));
```
2) BUG() inside biovec_slab(),
[code block 3]
```
   51         default:
   52                 BUG();
   53                 return NULL;
```

All the above panics are original from cached_dev_cache_miss() by the oversized parameter 'sectors'.

Inside cached_dev_cache_miss(), parameter 'sectors' is used to calculate the size of data read from backing device for the cache missing. This size is stored in s->insert_bio_sectors by the following lines of code,
[code block 4]
```
  909    s->insert_bio_sectors = min(sectors, bio_sectors(bio) + reada);
```

Then the actual key inserting to the internal B+ tree is generated and stored in s->iop.replace_key by the following lines of code,
[code block 5]
```
  911   s->iop.replace_key = KEY(s->iop.inode,
  912                     bio->bi_iter.bi_sector + s->insert_bio_sectors,
  913                     s->insert_bio_sectors);
```
The oversized parameter 'sectors' may trigger panic 1) by BUG_ON() from the above code block.

And the bio sending to backing device for the missing data is allocated with hint from s->insert_bio_sectors by the following lines of code,
[code block 6]
```
  926    cache_bio = bio_alloc_bioset(GFP_NOWAIT,
  927                 DIV_ROUND_UP(s->insert_bio_sectors, PAGE_SECTORS),
  928                 &dc->disk.bio_split);
```
The oversized parameter 'sectors' may trigger panic 2) by BUG() from the agove code block.

Now let me explain how the panics happen with the oversized 'sectors'. In code block 5, replace_key is generated by macro KEY(). From the definition of macro KEY(),
[code block 7]
```
  71 #define KEY(inode, offset, size)                                    \
  72 ((struct bkey) {                                                    \
  73      .high = (1ULL << 63) | ((__u64) (size) << 20) | (inode),     \
  74      .low = (offset)                                               \
  75 })
```

Here 'size' is 16bits width embedded in 64bits member 'high' of struct bkey. But in code block 1, if "KEY_START(k) - bio->bi_iter.bi_sector" is very probably to be larger than (1<<16) - 1, which makes the bkey size calculation in code block 5 is overflowed. In one bug report the value of parameter 'sectors' is 131072 (= 1 << 17), the overflowed 'sectors' results the overflowed s->insert_bio_sectors in code block 4, then makes size field of s->iop.replace_key to be 0 in code block 5. Then the 0-sized s->iop.replace_key is inserted into the internal B+ tree as cache missing check key (a special key to detect and avoid a racing between normal write request and cache missing read request) as,
[code block 8]
```
  915   ret = bch_btree_insert_check_key(b, &s->op, &s->iop.replace_key);
```

Then the 0-sized s->iop.replace_key as 3rd parameter triggers the bkey size check BUG_ON() in code block 2, and causes the kernel panic 1).

Another kernel panic is from code block 6, is by the bvecs number

```
    oversized value s->insert_bio_sectors from code block 4,
            min(sectors, bio_sectors(bio) + reada)
    There are two possibility for oversized reresult,
    - bio_sectors(bio) is valid, but bio_sectors(bio) + reada is oversized.
    - sectors < bio_sectors(bio) + reada, but sectors is oversized.

    From a bug report the result of "DIV_ROUND_UP(s->insert_bio_sectors,
    PAGE_SECTORS)" from code block 6 can be 344, 282, 946, 342 and many
    other values which larther than BIO_MAX_VECS (a.k.a 256). When calling
    bio_alloc_bioset() with such larger-than-256 value as the 2nd parameter,
    this value will eventually be sent to biovec_slab() as parameter
    'nr_vecs' in following code path,
       bio_alloc_bioset() ==> bvec_alloc() ==> biovec_slab()
    Because parameter 'nr_vecs' is larger-than-256 value, the panic by BUG()
    in code block 3 is triggered inside biovec_slab().

    From the above analysis, we know that the 4th parameter 'sector' sent
    into cached_dev_cache_miss() may cause overflow in code block 5 and 6,
    and finally cause kernel panic in code block 2 and 3. And if result of
    bio_sectors(bio) + reada exceeds valid bvecs number, it may also trigger
    kernel panic in code block 3 from code block 6.

    Now the almost-useless readahead size for cache missing request back to
    backing device is removed, this patch can fix the oversized issue with
    more simpler method.
    - add a local variable size_limit,  set it by the minimum value from
      the max bkey size and max bio bvecs number.
    - set s->insert_bio_sectors by the minimum value from size_limit,
      sectors, and the sectors size of bio.
    - replace sectors by s->insert_bio_sectors to do bio_next_split.

    By the above method with size_limit, s->insert_bio_sectors will never
    result oversized replace_key size or bio bvecs number. And split bio
    'miss' from bio_next_split() will always match the size of 'cache_bio',
    that is the current maximum bio size we can sent to backing device for
    fetching the cache missing data.

    Current problmatic code can be partially found since Linux v3.13-rc1,
    therefore all maintained stable kernels should try to apply this fix.

    Reported-by: Alexander Ullrich <ealex1979@gmail.com>
    Reported-by: Diego Ercolani <diego.ercolani@gmail.com>
    Reported-by: Jan Szubiak <jan.szubiak@linuxpolska.pl>
    Reported-by: Marco Rebhan <me@dblsaiko.net>
    Reported-by: Matthias Ferdinand <bcache@mfedv.net>
    Reported-by: Victor Westerhuis <victor@westerhu.is>
    Reported-by: Vojtech Pavlik <vojtech@suse.cz>
    Reported-and-tested-by: Rolf Fokkens <rolf@rolffokkens.nl>
    Reported-and-tested-by: Thorsten Knabe <linux@thorsten-knabe.de>
    Signed-off-by: Coly Li <colyli@suse.de>
    Cc: stable@vger.kernel.org
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Kent Overstreet <kent.overstreet@gmail.com>
    Cc: Nix <nix@esperi.org.uk>
    Cc: Takashi Iwai <tiwai@suse.com>
    Link: https://lore.kernel.org/r/20210607125052.21277-3-colyli@suse.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7550b63e30a5be89045d2a4cce9d84f3128ed2c9
Author: Coly Li <colyli@suse.de>
Date:   Mon Jun 7 20:50:51 2021 +0800

    bcache: remove bcache device self-defined readahead
```

```
    commit 1616a4c2ab1a80893b6890ae93da40a2b1d0c691 upstream.

    For read cache missing, bcache defines a readahead size for the read I/O
    request to the backing device for the missing data. This readahead size
    is initialized to 0, and almost no one uses it to avoid unnecessary read
    amplifying onto backing device and write amplifying onto cache device.
    Considering upper layer file system code has readahead logic allready
    and works fine with readahead_cache_policy sysfile interface, we don't
    have to keep bcache self-defined readahead anymore.

    This patch removes the bcache self-defined readahead for cache missing
    request for backing device, and the readahead sysfs file interfaces are
    removed as well.

    This is the preparation for next patch to fix potential kernel panic due
    to oversized request in a simpler method.

    Reported-by: Alexander Ullrich <ealex1979@gmail.com>
    Reported-by: Diego Ercolani <diego.ercolani@gmail.com>
    Reported-by: Jan Szubiak <jan.szubiak@linuxpolska.pl>
    Reported-by: Marco Rebhan <me@dblsaiko.net>
    Reported-by: Matthias Ferdinand <bcache@mfedv.net>
    Reported-by: Victor Westerhuis <victor@westerhu.is>
    Reported-by: Vojtech Pavlik <vojtech@suse.cz>
    Reported-and-tested-by: Rolf Fokkens <rolf@rolffokkens.nl>
    Reported-and-tested-by: Thorsten Knabe <linux@thorsten-knabe.de>
    Signed-off-by: Coly Li <colyli@suse.de>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Cc: stable@vger.kernel.org
    Cc: Kent Overstreet <kent.overstreet@gmail.com>
    Cc: Nix <nix@esperi.org.uk>
    Cc: Takashi Iwai <tiwai@suse.com>
    Link: https://lore.kernel.org/r/20210607125052.21277-2-colyli@suse.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 616126bff40c9d9afeabab0f48884a1c25b28c48
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Thu May 27 13:34:41 2021 +0200

    tick/nohz: Only check for RCU deferred wakeup on user/guest entry when needed

    commit f268c3737ecaefcfeecfb4cb5e44958a8976f067 upstream.

    Checking for and processing RCU-nocb deferred wakeup upon user/guest
    entry is only relevant when nohz_full runs on the local CPU, otherwise
    the periodic tick should take care of it.

    Make sure we don't needlessly pollute these fast-paths as a -3%
    performance regression on a will-it-scale.per_process_ops has been
    reported so far.

    Fixes: 47b8ff194c1f (entry: Explicitly flush pending rcuog wakeup before last
rescheduling point)
    Fixes: 4ae7dc97f726 (entry/kvm: Explicitly flush pending rcuog wakeup before last
rescheduling point)
    Reported-by: kernel test robot <oliver.sang@intel.com>
    Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Paul E. McKenney <paulmck@kernel.org>
    Cc: stable@vger.kernel.org
    Link: https://lkml.kernel.org/r/20210527113441.465489-1-frederic@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c597d5c59c7a6417dba06590f59b922e01188e8d
```

```
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Sun Jun 6 21:07:40 2021 +0530

    bus: mhi: pci_generic: Fix possible use-after-free in mhi_pci_remove()

    commit 0b67808ade8893a1b3608ddd74fac7854786c919 upstream.

    This driver's remove path calls del_timer(). However, that function
    does not wait until the timer handler finishes. This means that the
    timer handler may still be running after the driver's remove function
    has finished, which would result in a use-after-free.

    Fix by calling del_timer_sync(), which makes sure the timer handler
    has finished, and unable to re-schedule itself.

    Link: https://lore.kernel.org/r/20210413160318.2003699-1-weiyongjun1@huawei.com
    Fixes: 8562d4fe34a3 ("mhi: pci_generic: Add health-check")
    Cc: stable <stable@vger.kernel.org>
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Reviewed-by: Hemant kumar <hemantk@codeaurora.org>
    Reviewed-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Reviewed-by: Loic Poulain <loic.poulain@linaro.org>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20210606153741.20725-3-manivannan.sadhasivam@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0bb28bc0457de549716bb562b5fd6192ced0013a
Author: Ricky Wu <ricky_wu@realtek.com>
Date:   Mon Jun 7 18:16:34 2021 +0800

    misc: rtsx: separate aspm mode into MODE_REG and MODE_CFG

    commit 3df4fce739e2b263120f528c5e0fe6b2f8937b5b upstream.

    aspm (Active State Power Management)
    rtsx_comm_set_aspm: this function is for driver to make sure
    not enter power saving when processing of init and card_detcct
    ASPM_MODE_CFG: 8411 5209 5227 5229 5249 5250
    Change back to use original way to control aspm
    ASPM_MODE_REG: 5227A 524A 5250A 5260 5261 5228
    Keep the new way to control aspm

    Fixes: 121e9c6b5c4c ("misc: rtsx: modify and fix init_hw function")
    Reported-by: Chris Chiu <chris.chiu@canonical.com>
    Tested-by: Gordon Lack <gordon.lack@dsl.pipex.com>
    Cc: stable <stable@vger.kernel.org>
    Signed-off-by: Ricky Wu <ricky_wu@realtek.com>
    Link: https://lore.kernel.org/r/20210607101634.4948-1-ricky_wu@realtek.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2147684be1ebdaf845783139b9bc4eba3fecd9e4
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Sat Jun 5 16:09:14 2021 +0800

    usb: misc: brcmstb-usb-pinmap: check return value after calling platform_get_resource()

    commit fbf649cd6d64d40c03c5397ecd6b1ae922ba7afc upstream.

    It will cause null-ptr-deref if platform_get_resource() returns NULL,
    we need check the return value.

    Fixes: 517c4c44b323 ("usb: Add driver to allow any GPIO to be used for 7211 USB
signals")
    Cc: stable <stable@vger.kernel.org>
```

```
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20210605080914.2057758-1-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f44f1fc30dae491e03c3c282510d3657f7fac98f
Author: Wenli Looi <wlooi@ucalgary.ca>
Date:   Mon Jun 7 23:46:20 2021 -0700

    staging: rtl8723bs: Fix uninitialized variables

    commit 43c85d770db80cb135f576f8fde6ff1a08e707a4 upstream.

    The sinfo.pertid and sinfo.generation variables are not initialized and
    it causes a crash when we use this as a wireless access point.

    [  456.873025] ------------[ cut here ]------------
    [  456.878198] kernel BUG at mm/slub.c:3968!
    [  456.882680] Internal error: Oops - BUG: 0 [#1] PREEMPT SMP ARM

      [ snip ]

    [  457.271004] Backtrace:
    [  457.273733] [<c02b7ee4>] (kfree) from [<c0e2a470>] (nl80211_send_station+0x954/0xfc4)
    [  457.282481]  r9:eccca0c0 r8:e8edfec0 r7:00000000 r6:00000011 r5:e80a9480 r4:e8edfe00
    [  457.291132] [<c0e29b1c>] (nl80211_send_station) from [<c0e2b18c>]
(cfg80211_new_sta+0x90/0x1cc)
    [  457.300850]  r10:e80a9480 r9:e8edfe00 r8:ea678cca r7:00000a20 r6:00000000 r5:ec46d000
    [  457.309586]  r4:ec46d9e0
    [  457.312433] [<c0e2b0fc>] (cfg80211_new_sta) from [<bf086684>]
(rtw_cfg80211_indicate_sta_assoc+0x80/0x9c [r8723bs])
    [  457.324095]  r10:00009930 r9:e85b9d80 r8:bf091050 r7:00000000 r6:00000000 r5:0000001c
    [  457.332831]  r4:c1606788
    [  457.335692] [<bf086604>] (rtw_cfg80211_indicate_sta_assoc [r8723bs]) from
[<bf03df38>] (rtw_stassoc_event_callback+0x1c8/0x1d4 [r8723bs])
    [  457.349489]  r7:ea678cc0 r6:000000a1 r5:f1225f84 r4:f086b000
    [  457.355845] [<bf03dd70>] (rtw_stassoc_event_callback [r8723bs]) from [<bf048e4c>]
(mlme_evt_hdl+0x8c/0xb4 [r8723bs])
    [  457.367601]  r7:c1604900 r6:f086c4b8 r5:00000000 r4:f086c000
    [  457.373959] [<bf048dc0>] (mlme_evt_hdl [r8723bs]) from [<bf03693c>]
(rtw_cmd_thread+0x198/0x3d8 [r8723bs])
    [  457.384744]  r5:f086e000 r4:f086c000
    [  457.388754] [<bf0367a4>] (rtw_cmd_thread [r8723bs]) from [<c014a214>]
(kthread+0x170/0x174)
    [  457.398083]  r10:ed7a57e8 r9:bf0367a4 r8:f086b000 r7:e8ede000 r6:00000000 r5:e9975200
    [  457.406828]  r4:e8369900
    [  457.409653] [<c014a0a4>] (kthread) from [<c01010e8>] (ret_from_fork+0x14/0x2c)
    [  457.417718] Exception stack(0xe8edffb0 to 0xe8edfff8)
    [  457.423356] ffa0:                                     00000000 00000000 00000000
00000000
    [  457.432492] ffc0: 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000
    [  457.441618] ffe0: 00000000 00000000 00000000 00000000 00000013 00000000
    [  457.449006]  r10:00000000 r9:00000000 r8:00000000 r7:00000000 r6:00000000 r5:c014a0a4
    [  457.457750]  r4:e9975200
    [  457.460574] Code: 1a000003 e5953004 e3130001 1a000000 (e7f001f2)
    [  457.467381] ---[ end trace 4acbc8c15e9e6aa7 ]---

    Link: https://forum.armbian.com/topic/14727-wifi-ap-kernel-bug-in-kernel-5444/
    Fixes: 8689c051a201 ("cfg80211: dynamically allocate per-tid stats for station info")
    Fixes: f5ea9120be2e ("nl80211: add generation number to all dumps")
    Signed-off-by: Wenli Looi <wlooi@ucalgary.ca>
    Reviewed-by: Dan Carpenter <dan.carpenter@oracle.com>
    Cc: stable <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210608064620.74059-1-wlooi@ucalgary.ca
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit ed0e2a893092c7fcb4ff7ba74e5efce53a6f5940
Author: Paolo Bonzini <pbonzini@redhat.com>
Date:   Tue Jun 8 15:31:42 2021 -0400

    kvm: avoid speculation-based attacks from out-of-range memslot accesses

    commit da27a83fd6cc7780fea190e1f5c19e87019da65c upstream.

    KVM's mechanism for accessing guest memory translates a guest physical
    address (gpa) to a host virtual address using the right-shifted gpa
    (also known as gfn) and a struct kvm_memory_slot.  The translation is
    performed in __gfn_to_hva_memslot using the following formula:

          hva = slot->userspace_addr + (gfn - slot->base_gfn) * PAGE_SIZE

    It is expected that gfn falls within the boundaries of the guest's
    physical memory.  However, a guest can access invalid physical addresses
    in such a way that the gfn is invalid.

    __gfn_to_hva_memslot is called from kvm_vcpu_gfn_to_hva_prot, which first
    retrieves a memslot through __gfn_to_memslot.  While __gfn_to_memslot
    does check that the gfn falls within the boundaries of the guest's
    physical memory or not, a CPU can speculate the result of the check and
    continue execution speculatively using an illegal gfn. The speculation
    can result in calculating an out-of-bounds hva.  If the resulting host
    virtual address is used to load another guest physical address, this
    is effectively a Spectre gadget consisting of two consecutive reads,
    the second of which is data dependent on the first.

    Right now it's not clear if there are any cases in which this is
    exploitable.  One interesting case was reported by the original author
    of this patch, and involves visiting guest page tables on x86.  Right
    now these are not vulnerable because the hva read goes through get_user(),
    which contains an LFENCE speculation barrier.  However, there are
    patches in progress for x86 uaccess.h to mask kernel addresses instead of
    using LFENCE; once these land, a guest could use speculation to read
    from the VMM's ring 3 address space.  Other architectures such as ARM
    already use the address masking method, and would be susceptible to
    this same kind of data-dependent access gadgets.  Therefore, this patch
    proactively protects from these attacks by masking out-of-bounds gfns
    in __gfn_to_hva_memslot, which blocks speculation of invalid hvas.

    Sean Christopherson noted that this patch does not cover
    kvm_read_guest_offset_cached.  This however is limited to a few bytes
    past the end of the cache, and therefore it is unlikely to be useful in
    the context of building a chain of data dependent accesses.

    Reported-by: Artemiy Margaritov <artemiy.margaritov@gmail.com>
    Co-developed-by: Artemiy Margaritov <artemiy.margaritov@gmail.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7707f7244d813f8279143d5dfdee7a318801231b
Author: Lai Jiangshan <laijs@linux.alibaba.com>
Date:   Thu Jun 3 13:24:55 2021 +0800

    KVM: X86: MMU: Use the correct inherited permissions to get shadow page

    commit b1bd5cba3306691c771d558e94baa73e8b0b96b7 upstream.

    When computing the access permissions of a shadow page, use the effective
    permissions of the walk up to that point, i.e. the logic AND of its parents'
    permissions.  Two guest PxE entries that point at the same table gfn need to
    be shadowed with different shadow pages if their parents' permissions are
```

```
    different.  KVM currently uses the effective permissions of the last
    non-leaf entry for all non-leaf entries.  Because all non-leaf SPTEs have
    full ("uwx") permissions, and the effective permissions are recorded only
    in role.access and merged into the leaves, this can lead to incorrect
    reuse of a shadow page and eventually to a missing guest protection page
    fault.

    For example, here is a shared pagetable:

       pgd[]   pud[]        pmd[]            virtual address pointers
                         /->pmd1(u--)->pte1(uw-)->page1 <- ptr1 (u--)
            /->pud1(uw-)--->pmd2(uw-)->pte2(uw-)->page2 <- ptr2 (uw-)
       pgd-|           (shared pmd[] as above)
            \->pud2(u--)--->pmd1(u--)->pte1(uw-)->page1 <- ptr3 (u--)
                         \->pmd2(uw-)->pte2(uw-)->page2 <- ptr4 (u--)

      pud1 and pud2 point to the same pmd table, so:
      - ptr1 and ptr3 points to the same page.
      - ptr2 and ptr4 points to the same page.

    (pud1 and pud2 here are pud entries, while pmd1 and pmd2 here are pmd entries)

    - First, the guest reads from ptr1 first and KVM prepares a shadow
      page table with role.access=u--, from ptr1's pud1 and ptr1's pmd1.
      "u--" comes from the effective permissions of pgd, pud1 and
      pmd1, which are stored in pt->access.  "u--" is used also to get
      the pagetable for pud1, instead of "uw-".

    - Then the guest writes to ptr2 and KVM reuses pud1 which is present.
      The hypervisor set up a shadow page for ptr2 with pt->access is "uw-"
      even though the pud1 pmd (because of the incorrect argument to
      kvm_mmu_get_page in the previous step) has role.access="u--".

    - Then the guest reads from ptr3.  The hypervisor reuses pud1's
      shadow pmd for pud2, because both use "u--" for their permissions.
      Thus, the shadow pmd already includes entries for both pmd1 and pmd2.

    - At last, the guest writes to ptr4.  This causes no vmexit or pagefault,
      because pud1's shadow page structures included an "uw-" page even though
      its role.access was "u--".

    Any kind of shared pagetable might have the similar problem when in
    virtual machine without TDP enabled if the permissions are different
    from different ancestors.

    In order to fix the problem, we change pt->access to be an array, and
    any access in it will not include permissions ANDed from child ptes.

    The test code is: https://lore.kernel.org/kvm/20210603050537.19605-1-
jiangshanlai@gmail.com/
    Remember to test it with TDP disabled.

    The problem had existed long before the commit 41074d07c78b ("KVM: MMU:
    Fix inherited permissions for emulated guest pte updates"), and it
    is hard to find which is the culprit.  So there is no fixes tag here.

    Signed-off-by: Lai Jiangshan <laijs@linux.alibaba.com>
    Message-Id: <20210603052455.21023-1-jiangshanlai@gmail.com>
    Cc: stable@vger.kernel.org
    Fixes: cea0f0e7ea54 ("[PATCH] KVM: MMU: Shadow page table caching")
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1f297c56bfc6183f9bf7952e3256ddd2a3d9e431
Author: Lai Jiangshan <laijs@linux.alibaba.com>
```

```
Date:   Tue Jun 1 01:22:56 2021 +0800

    KVM: x86: Unload MMU on guest TLB flush if TDP disabled to force MMU sync

    commit b53e84eed08b88fd3ff59e5c2a7f1a69d4004e32 upstream.

    When using shadow paging, unload the guest MMU when emulating a guest TLB
    flush to ensure all roots are synchronized.  From the guest's perspective,
    flushing the TLB ensures any and all modifications to its PTEs will be
    recognized by the CPU.

    Note, unloading the MMU is overkill, but is done to mirror KVM's existing
    handling of INVPCID(all) and ensure the bug is squashed.  Future cleanup
    can be done to more precisely synchronize roots when servicing a guest
    TLB flush.

    If TDP is enabled, synchronizing the MMU is unnecessary even if nested
    TDP is in play, as a "legacy" TLB flush from L1 does not invalidate L1's
    TDP mappings.  For EPT, an explicit INVEPT is required to invalidate
    guest-physical mappings; for NPT, guest mappings are always tagged with
    an ASID and thus can only be invalidated via the VMCB's ASID control.

    This bug has existed since the introduction of KVM_VCPU_FLUSH_TLB.
    It was only recently exposed after Linux guests stopped flushing the
    local CPU's TLB prior to flushing remote TLBs (see commit 4ce94eabac16,
    "x86/mm/tlb: Flush remote and local TLBs concurrently"), but is also
    visible in Windows 10 guests.

    Tested-by: Maxim Levitsky <mlevitsk@redhat.com>
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    Fixes: f38a7b75267f ("KVM: X86: support paravirtualized help for TLB shootdowns")
    Signed-off-by: Lai Jiangshan <laijs@linux.alibaba.com>
    [sean: massaged comment and changelog]
    Message-Id: <20210531172256.2908-1-jiangshanlai@gmail.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 82cc1c4728ce3a72829d2a4d2b42e57739d4c438
Author: Kan Liang <kan.liang@linux.intel.com>
Date:   Tue Jun 1 06:09:03 2021 -0700

    perf/x86/intel/uncore: Fix M2M event umask for Ice Lake server

    commit 848ff3768684701a4ce73a2ec0e5d438d4e2b0da upstream.

    Perf tool errors out with the latest event list for the Ice Lake server.

    event syntax error: 'unc_m2m_imc_reads.to_pmm'
                                \___ value too big for format, maximum is 255

    The same as the Snow Ridge server, the M2M uncore unit in the Ice Lake
    server has the unit mask extension field as well.

    Fixes: 2b3b76b5ec67 ("perf/x86/intel/uncore: Add Ice Lake server uncore support")
    Reported-by: Jin Yao <yao.jin@linux.intel.com>
    Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Cc: stable@vger.kernel.org
    Link: https://lkml.kernel.org/r/1622552943-119174-1-git-send-email-
kan.liang@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8c9987ad71341731f457adeef59bb1b8336742a2
```

```
Author: Tor Vic <torvic9@mailbox.org>
Date:   Thu Jun 10 20:58:06 2021 +0000

    x86, lto: Pass -stack-alignment only on LLD < 13.0.0

    commit 2398ce80152aae33b9501ef54452e09e8e8d4262 upstream.

    Since LLVM commit 3787ee4, the '-stack-alignment' flag has been dropped
    [1], leading to the following error message when building a LTO kernel
    with Clang-13 and LLD-13:

        ld.lld: error: -plugin-opt=-: ld.lld: Unknown command line argument
        '-stack-alignment=8'.  Try 'ld.lld --help'
        ld.lld: Did you mean '--stackrealign=8'?

    It also appears that the '-code-model' flag is not necessary anymore
    starting with LLVM-9 [2].

    Drop '-code-model' and make '-stack-alignment' conditional on LLD < 13.0.0.

    These flags were necessary because these flags were not encoded in the
    IR properly, so the link would restart optimizations without them. Now
    there are properly encoded in the IR, and these flags exposing
    implementation details are no longer necessary.

    [1] https://reviews.llvm.org/D103048
    [2] https://reviews.llvm.org/D52322

    Cc: stable@vger.kernel.org
    Link: https://github.com/ClangBuiltLinux/linux/issues/1377
    Signed-off-by: Tor Vic <torvic9@mailbox.org>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Tested-by: Nathan Chancellor <nathan@kernel.org>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/f2c018ee-5999-741e-58d4-e482d5246067@mailbox.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit dfedcac75b81e617b3a554c76e817d7961637a33
Author: Desmond Cheong Zhi Xi <desmondcheongzx@gmail.com>
Date:   Wed Jun 9 17:21:19 2021 +0800

    drm: Lock pointer access in drm_master_release()

    commit c336a5ee984708db4826ef9e47d184e638e29717 upstream.

    This patch eliminates the following smatch warning:
    drivers/gpu/drm/drm_auth.c:320 drm_master_release() warn: unlocked access 'master' (line
318) expected lock '&dev->master_mutex'

    The 'file_priv->master' field should be protected by the mutex lock to
    '&dev->master_mutex'. This is because other processes can concurrently
    modify this field and free the current 'file_priv->master'
    pointer. This could result in a use-after-free error when 'master' is
    dereferenced in subsequent function calls to
    'drm_legacy_lock_master_cleanup()' or to 'drm_lease_revoke()'.

    An example of a scenario that would produce this error can be seen
    from a similar bug in 'drm_getunique()' that was reported by Syzbot:
    https://syzkaller.appspot.com/bug?id=148d2f1dfac64af52ffd27b661981a540724f803

    In the Syzbot report, another process concurrently acquired the
    device's master mutex in 'drm_setmaster_ioctl()', then overwrote
    'fpriv->master' in 'drm_new_set_master()'. The old value of
    'fpriv->master' was subsequently freed before the mutex was unlocked.
```

```
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Desmond Cheong Zhi Xi <desmondcheongzx@gmail.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210609092119.173590-1-
desmondcheongzx@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f773f8cccac13c7e7bbd9182e7996c727742488e
Author: Desmond Cheong Zhi Xi <desmondcheongzx@gmail.com>
Date:   Tue Jun 8 19:04:36 2021 +0800

    drm: Fix use-after-free read in drm_getunique()

    commit b436acd1cf7fac0ba987abd22955d98025c80c2b upstream.

    There is a time-of-check-to-time-of-use error in drm_getunique() due
    to retrieving file_priv->master prior to locking the device's master
    mutex.

    An example can be seen in the crash report of the use-after-free error
    found by Syzbot:
    https://syzkaller.appspot.com/bug?id=148d2f1dfac64af52ffd27b661981a540724f803

    In the report, the master pointer was used after being freed. This is
    because another process had acquired the device's master mutex in
    drm_setmaster_ioctl(), then overwrote fpriv->master in
    drm_new_set_master(). The old value of fpriv->master was subsequently
    freed before the mutex was unlocked.

    To fix this, we lock the device's master mutex before retrieving the
    pointer from from fpriv->master. This patch passes the Syzbot
    reproducer test.

    Reported-by: syzbot+c3a706cec1ea99e1c693@syzkaller.appspotmail.com
    Signed-off-by: Desmond Cheong Zhi Xi <desmondcheongzx@gmail.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210608110436.239583-1-
desmondcheongzx@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 135076c45e4561cd388cc21f30ee8bd1076b7de1
Author: Mika Westerberg <mika.westerberg@linux.intel.com>
Date:   Wed Jun 9 13:25:33 2021 +0300

    ACPI: Pass the same capabilities to the _OSC regardless of the query flag

    commit 159d8c274fd92438ca6d7068d7a5eeda157227f4 upstream.

    Commit 719e1f561afb ("ACPI: Execute platform _OSC also with query bit
    clear") makes acpi_bus_osc_negotiate_platform_control() not only query
    the platforms capabilities but it also commits the result back to the
    firmware to report which capabilities are supported by the OS back to
    the firmware

    On certain systems the BIOS loads SSDT tables dynamically based on the
    capabilities the OS claims to support. However, on these systems the
    _OSC actually clears some of the bits (under certain conditions) so what
    happens is that now when we call the _OSC twice the second time we pass
    the cleared values and that results errors like below to appear on the
    system log:

      ACPI BIOS Error (bug): Could not resolve symbol [\_PR.PR00._CPC], AE_NOT_FOUND
(20210105/psargs-330)
```

```
       ACPI Error: Aborting method \_PR.PR01._CPC due to previous error (AE_NOT_FOUND)
  (20210105/psparse-529)

      In addition the ACPI 6.4 spec says following [1]:

        If the OS declares support of a feature in the Support Field in one
        call to _OSC, then it must preserve the set state of that bit
        (declaring support for that feature) in all subsequent calls.

      Based on the above we can fix the issue by passing the same set of
      capabilities to the platform wide _OSC in both calls regardless of the
      query flag.

      While there drop the context.ret.length checks which were wrong to begin
      with (as the length is number of bytes not elements). This is already
      checked in acpi_run_osc() that also returns an error in that case.

      Includes fixes by Hans de Goede.

      [1]
  https://uefi.org/specs/ACPI/6.4/06_Device_Configuration/Device_Configuration.html#sequence-
  of-osc-calls

      BugLink: https://bugzilla.kernel.org/show_bug.cgi?id=213023
      BugLink: https://bugzilla.redhat.com/show_bug.cgi?id=1963717
      Fixes: 719e1f561afb ("ACPI: Execute platform _OSC also with query bit clear")
      Cc: 5.12+ <stable@vger.kernel.org> # 5.12+
      Signed-off-by: Mika Westerberg <mika.westerberg@linux.intel.com>
      Reviewed-by: Hans de Goede <hdegoede@redhat.com>
      Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 6192af07e37b2b2ab39215af9eb171d4e2e44a26
  Author: Zhang Rui <rui.zhang@intel.com>
  Date:   Thu Jun 3 10:34:14 2021 +0800

      Revert "ACPI: sleep: Put the FACS table after using it"

      commit f1ffa9d4cccc8fdf6c03fb1b3429154d22037988 upstream.

      Commit 95722237cb2a ("ACPI: sleep: Put the FACS table after using it")
      puts the FACS table during initialization.

      But the hardware signature bits in the FACS table need to be accessed,
      after every hibernation, to compare with the original hardware
      signature.

      So there is no reason to release the FACS table mapping after
      initialization.

      This reverts commit 95722237cb2ae4f7b73471058cdb19e8f4057c93.

      An alternative solution is to use acpi_gbl_FACS variable instead, which
      is mapped by the ACPICA core and never released.

      Link: https://bugzilla.kernel.org/show_bug.cgi?id=212277
      Reported-by: Stephan Hohe <sth.dev@tejp.de>
      Signed-off-by: Zhang Rui <rui.zhang@intel.com>
      Cc: 5.8+ <stable@vger.kernel.org> # 5.8+
      Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 01415ff85a24308059e06ca3e97fd7bf75648690
  Author: Lukas Wunner <lukas@wunner.de>
  Date:   Sat May 22 19:49:50 2021 +0200
```

```
    spi: bcm2835: Fix out-of-bounds access with more than 4 slaves

    commit 13817d466eb8713a1ffd254f537402f091d48444 upstream.

    Commit 571e31fa60b3 ("spi: bcm2835: Cache CS register value for
    ->prepare_message()") limited the number of slaves to 3 at compile-time.
    The limitation was necessitated by a statically-sized array prepare_cs[]
    in the driver private data which contains a per-slave register value.

    The commit sought to enforce the limitation at run-time by setting the
    controller's num_chipselect to 3:  Slaves with a higher chipselect are
    rejected by spi_add_device().

    However the commit neglected that num_chipselect only limits the number
    of *native* chipselects.  If GPIO chipselects are specified in the
    device tree for more than 3 slaves, num_chipselect is silently raised by
    of_spi_get_gpio_numbers() and the result are out-of-bounds accesses to
    the statically-sized array prepare_cs[].

    As a bandaid fix which is backportable to stable, raise the number of
    allowed slaves to 24 (which "ought to be enough for anybody"), enforce
    the limitation on slave ->setup and revert num_chipselect to 3 (which is
    the number of native chipselects supported by the controller).
    An upcoming for-next commit will allow an arbitrary number of slaves.

    Fixes: 571e31fa60b3 ("spi: bcm2835: Cache CS register value for ->prepare_message()")
    Reported-by: Joe Burmeister <joe.burmeister@devtank.co.uk>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Cc: stable@vger.kernel.org # v5.4+
    Cc: Phil Elwell <phil@raspberrypi.com>
    Link:
https://lore.kernel.org/r/75854affc1923309fde05e47494263bde73e5592.1621703210.git.lukas@wunn
er.de
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b7886783b01b5637a42a8dd813152e724c982804
Author: Jeremy Szu <jeremy.szu@canonical.com>
Date:   Tue Jun 8 19:47:48 2021 +0800

    ALSA: hda/realtek: fix mute/micmute LEDs for HP ZBook Power G8

    commit 600dd2a7e8b62170d177381cc1303861f48f9780 upstream.

    The HP ZBook Power G8 using ALC236 codec which using 0x02 to
    control mute LED and 0x01 to control micmute LED.
    Therefore, add a quirk to make it works.

    Signed-off-by: Jeremy Szu <jeremy.szu@canonical.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210608114750.32009-1-jeremy.szu@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e0173c323946f4a9c372eae0625f3a80e3e55b6f
Author: Jeremy Szu <jeremy.szu@canonical.com>
Date:   Sat Jun 5 16:25:38 2021 +0800

    ALSA: hda/realtek: fix mute/micmute LEDs for HP EliteBook 840 Aero G8

    commit dfb06401b4cdfc71e2fc3e19b877ab845cc9f7f7 upstream.

    The HP EliteBook 840 Aero G8 using ALC285 codec which using 0x04 to
    control mute LED and 0x01 to control micmute LED.
```

```
    In the other hand, there is no output from right channel of speaker.
    Therefore, add a quirk to make it works.

    Signed-off-by: Jeremy Szu <jeremy.szu@canonical.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210605082539.41797-3-jeremy.szu@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3513eda896a1b13473f997617fb99c1026d68400
Author: Jeremy Szu <jeremy.szu@canonical.com>
Date:   Sat Jun 5 16:25:37 2021 +0800

    ALSA: hda/realtek: fix mute/micmute LEDs and speaker for HP EliteBook x360 1040 G8

    commit 61d3e87468fad82dc8e8cb6de7db563ada64b532 upstream.

    The HP EliteBook x360 1040 G8 using ALC285 codec which using 0x04 to control
    mute LED and 0x01 to control micmute LED.
    In the other hand, there is no output from right channel of speaker.
    Therefore, add a quirk to make it works.

    Signed-off-by: Jeremy Szu <jeremy.szu@canonical.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210605082539.41797-2-jeremy.szu@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b5a06c7723411ac936ed1ebecf26d967c41dbf87
Author: Jeremy Szu <jeremy.szu@canonical.com>
Date:   Sat Jun 5 16:25:36 2021 +0800

    ALSA: hda/realtek: fix mute/micmute LEDs and speaker for HP Elite Dragonfly G2

    commit 15d295b560e6dd45f839a53ae69e4f63b54eb32f upstream.

    The HP Elite Dragonfly G2 using ALC285 codec which using 0x04 to control
    mute LED and 0x01 to control micmute LED.
    In the other hand, there is no output from right channel of speaker.
    Therefore, add a quirk to make it works.

    Signed-off-by: Jeremy Szu <jeremy.szu@canonical.com>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20210605082539.41797-1-jeremy.szu@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2901ac6275222eb90dac2a07da0e3092c7d62ac
Author: Hui Wang <hui.wang@canonical.com>
Date:   Tue Jun 8 10:46:00 2021 +0800

    ALSA: hda/realtek: headphone and mic don't work on an Acer laptop

    commit 57c9e21a49b1c196cda28f54de9a5d556ac93f20 upstream.

    There are 2 issues on this machine, the 1st one is mic's plug/unplug
    can't be detected, that is because the mic is set to manual detecting
    mode, need to apply ALC255_FIXUP_XIAOMI_HEADSET_MIC to set it to auto
    detecting mode. The other one is headphone's plug/unplug can't be
    detected by pulseaudio, that is because the pulseaudio will use
    ucm2/sof-hda-dsp on this machine, and the ucm2 only handle
    'Headphone Jack', but on this machine the headphone's pincfg sets the
    location to Front, then the alsa mixer name is "Front Headphone Jack"
    instead of "Headphone Jack", so override the pincfg to change location
    to Left.
```

```
    BugLink: http://bugs.launchpad.net/bugs/1930188
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Hui Wang <hui.wang@canonical.com>
    Link: https://lore.kernel.org/r/20210608024600.6198-1-hui.wang@canonical.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 05e4deffafdaefb46e201b2cbf300abf310e6292
Author: Takashi Sakamoto <o-takashi@sakamocchi.jp>
Date:   Sat Jun 5 18:10:54 2021 +0900

    ALSA: firewire-lib: fix the context to call snd_pcm_stop_xrun()

    commit 9981b20a5e3694f4625ab5a1ddc98ce7503f6d12 upstream.

    In the workqueue to queue wake-up event, isochronous context is not
    processed, thus it's useless to check context for the workqueue to switch
    status of runtime for PCM substream to XRUN. On the other hand, in
    software IRQ context of 1394 OHCI, it's needed.

    This commit fixes the bug introduced when tasklet was replaced with
    workqueue.

    Cc: <stable@vger.kernel.org>
    Fixes: 2b3d2987d800 ("ALSA: firewire: Replace tasklet with work")
    Signed-off-by: Takashi Sakamoto <o-takashi@sakamocchi.jp>
    Link: https://lore.kernel.org/r/20210605091054.68866-1-o-takashi@sakamocchi.jp
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 536a7646c00a0f14fee49e5e313109e5da2f6031
Author: Takashi Iwai <tiwai@suse.de>
Date:   Thu Jun 10 17:20:59 2021 +0200

    ALSA: seq: Fix race of snd_seq_timer_open()

    commit 83e197a8414c0ba545e7e3916ce05f836f349273 upstream.

    The timer instance per queue is exclusive, and snd_seq_timer_open()
    should have managed the concurrent accesses.  It looks as if it's
    checking the already existing timer instance at the beginning, but
    it's not right, because there is no protection, hence any later
    concurrent call of snd_seq_timer_open() may override the timer
    instance easily.  This may result in UAF, as the leftover timer
    instance can keep running while the queue itself gets closed, as
    spotted by syzkaller recently.

    For avoiding the race, add a proper check at the assignment of
    tmr->timeri again, and return -EBUSY if it's been already registered.

    Reported-by: syzbot+ddc1260a83ed1cbf6fb5@syzkaller.appspotmail.com
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/000000000000dce34f05c42f110c@google.com
    Link: https://lore.kernel.org/r/20210610152059.24633-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9f573a752d1b3180f78d473e8cf0c4bda2ea30a9
Author: Chris Packham <chris.packham@alliedtelesis.co.nz>
Date:   Wed May 12 09:20:52 2021 +1200

    i2c: mpc: implement erratum A-004447 workaround

    [ Upstream commit drivers/i2c/busses/i2c-mpc.c ]
```

```
    The P2040/P2041 has an erratum where the normal i2c recovery mechanism
    does not work. Implement the alternative recovery mechanism documented
    in the P2040 Chip Errata Rev Q.

    Signed-off-by: Chris Packham <chris.packham@alliedtelesis.co.nz>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a08e6771f30c74cef01dfd3d6821db6b17b4a199
Author: Chris Packham <chris.packham@alliedtelesis.co.nz>
Date:   Mon Mar 29 14:52:03 2021 +1300

    i2c: mpc: Make use of i2c_recover_bus()

    [ Upstream commit 65171b2df15eb7545431d75c2729b5062da89b43 ]

    Move the existing calls of mpc_i2c_fixup() to a recovery function
    registered via bus_recovery_info. This makes it more obvious that
    recovery is supported and allows for a future where recovery is
    triggered by the i2c core.

    Signed-off-by: Chris Packham <chris.packham@alliedtelesis.co.nz>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1cbb5b921e810d84bd52817148c0905260a08189
Author: Lukas Wunner <lukas@wunner.de>
Date:   Thu May 27 23:10:56 2021 +0200

    spi: Cleanup on failure of initial setup

    [ Upstream commit 2ec6f20b33eb4f62ab90bdcd620436c883ec3af6 ]

    Commit c7299fea6769 ("spi: Fix spi device unregister flow") changed the
    SPI core's behavior if the ->setup() hook returns an error upon adding
    an spi_device:  Before, the ->cleanup() hook was invoked to free any
    allocations that were made by ->setup().  With the commit, that's no
    longer the case, so the ->setup() hook is expected to free the
    allocations itself.

    I've identified 5 drivers which depend on the old behavior and am fixing
    them up hereinafter: spi-bitbang.c spi-fsl-spi.c spi-omap-uwire.c
    spi-omap2-mcspi.c spi-pxa2xx.c

    Importantly, ->setup() is not only invoked on spi_device *addition*:
    It may subsequently be called to *change* SPI parameters.  If changing
    these SPI parameters fails, freeing memory allocations would be wrong.
    That should only be done if the spi_device is finally destroyed.
    I am therefore using a bool "initial_setup" in 4 of the affected drivers
    to differentiate between the invocation on *adding* the spi_device and
    any subsequent invocations: spi-bitbang.c spi-fsl-spi.c spi-omap-uwire.c
    spi-omap2-mcspi.c

    In spi-pxa2xx.c, it seems the ->setup() hook can only fail on spi_device
    addition, not any subsequent calls.  It therefore doesn't need the bool.

    It's worth noting that 5 other drivers already perform a cleanup if the
    ->setup() hook fails.  Before c7299fea6769, they caused a double-free
    if ->setup() failed on spi_device addition.  Since the commit, they're
    fine.  These drivers are: spi-mpc512x-psc.c spi-pl022.c spi-s3c64xx.c
    spi-st-ssc4.c spi-tegra114.c

    (spi-pxa2xx.c also already performs a cleanup, but only in one of
    several error paths.)
```

```
    Fixes: c7299fea6769 ("spi: Fix spi device unregister flow")
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Cc: Saravana Kannan <saravanak@google.com>
    Acked-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com> # pxa2xx
    Link:
https://lore.kernel.org/r/f76a0599469f265b69c371538794101fa37b5536.1622149321.git.lukas@wunn
er.de
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3576a4c1f8371100e533bb31c419ed0598d1c8d
Author: Saravana Kannan <saravanak@google.com>
Date:   Wed May 5 09:47:34 2021 -0700

    spi: Don't have controller clean up spi device before driver unbind

    [ Upstream commit 27e7db56cf3dffd302bd7ddfacb1d405cf671a2a ]

    When a spi device is unregistered and triggers a driver unbind, the
    driver might need to access the spi device. So, don't have the
    controller clean up the spi device before the driver is unbound. Clean
    up the spi device after the driver is unbound.

    Fixes: c7299fea6769 ("spi: Fix spi device unregister flow")
    Reported-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Tested-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Link: https://lore.kernel.org/r/20210505164734.175546-1-saravanak@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d5773d335adf7ccf8982f5987d77d01cd02cf75
Author: Chris Packham <chris.packham@alliedtelesis.co.nz>
Date:   Wed May 12 09:20:51 2021 +1200

    powerpc/fsl: set fsl,i2c-erratum-a004447 flag for P1010 i2c controllers

    [ Upstream commit 19ae697a1e4edf1d755b413e3aa38da65e2db23b ]

    The i2c controllers on the P1010 have an erratum where the documented
    scheme for i2c bus recovery will not work (A-004447). A different
    mechanism is needed which is documented in the P1010 Chip Errata Rev L.

    Signed-off-by: Chris Packham <chris.packham@alliedtelesis.co.nz>
    Acked-by: Michael Ellerman <mpe@ellerman.id.au>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a8089545a34f696e04e710d7fcccc95eb2e03180
Author: Chris Packham <chris.packham@alliedtelesis.co.nz>
Date:   Wed May 12 09:20:50 2021 +1200

    powerpc/fsl: set fsl,i2c-erratum-a004447 flag for P2041 i2c controllers

    [ Upstream commit 7adc7b225cddcfd0f346d10144fd7a3d3d9f9ea7 ]

    The i2c controllers on the P2040/P2041 have an erratum where the
    documented scheme for i2c bus recovery will not work (A-004447). A
    different mechanism is needed which is documented in the P2040 Chip
    Errata Rev Q (latest available at the time of writing).

    Signed-off-by: Chris Packham <chris.packham@alliedtelesis.co.nz>
    Acked-by: Michael Ellerman <mpe@ellerman.id.au>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
```

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

commit 983b52636a439368272313585ab27d7addc92bf8
Author: Sagi Grimberg <sagi@grimberg.me>
Date:   Tue May 25 08:49:05 2021 -0700

    nvmet: fix false keep-alive timeout when a controller is torn down

    [ Upstream commit aaeadd7075dc9e184bc7876e9dd7b3bada771df2 ]

    Controller teardown flow may take some time in case it has many I/O
    queues, and the host may not send us keep-alive during this period.
    Hence reset the traffic based keep-alive timer so we don't trigger
    a controller teardown as a result of a keep-alive expiration.

    Reported-by: Yi Zhang <yi.zhang@redhat.com>
    Signed-off-by: Sagi Grimberg <sagi@grimberg.me>
    Reviewed-by: Chaitanya Kulkarni <chaitanya.kulkarni@wdc.com>
    Reviewed-by: Hannes Reinecke <hare@suse.de>
    Tested-by: Yi Zhang <yi.zhang@redhat.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c4082a91c49f787ee13f977f47717cfd29ec5519
Author: Sagi Grimberg <sagi@grimberg.me>
Date:   Fri May 21 14:51:15 2021 -0700

    nvme-tcp: remove incorrect Kconfig dep in BLK_DEV_NVME

    [ Upstream commit 042a3eaad6daeabcfaf163aa44da8ea3cf8b5496 ]

    We need to select NVME_CORE.

    Signed-off-by: Sagi Grimberg <sagi@grimberg.me>
    Reviewed-by: Max Gurtovoy <mgurtovoy@nvidia.com>
    Reviewed-by: Chaitanya Kulkarni <chaitanya.kulkarni@wdc.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8086763d3759a9fd3618927fc60f42b25bbacb5d
Author: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
Date:   Tue May 25 19:00:12 2021 +0800

    bnx2x: Fix missing error code in bnx2x_iov_init_one()

    [ Upstream commit 65161c35554f7135e6656b3df1ce2c500ca0bdcf ]

    Eliminate the follow smatch warning:

    drivers/net/ethernet/broadcom/bnx2x/bnx2x_sriov.c:1227
    bnx2x_iov_init_one() warn: missing error code 'err'.

    Reported-by: Abaci Robot <abaci@linux.alibaba.com>
    Signed-off-by: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fb429fad145a987056ac315b2c859197c3daa890
Author: John Keeping <john@metanate.com>
Date:   Wed May 12 12:14:21 2021 +0100

    dm verity: fix require_signatures module_param permissions

    [ Upstream commit 0c1f3193b1cdd21e7182f97dc9bca7d284d18a15 ]
```

```
    The third parameter of module_param() is permissions for the sysfs node
    but it looks like it is being used as the initial value of the parameter
    here.  In fact, false here equates to omitting the file from sysfs and
    does not affect the value of require_signatures.

    Making the parameter writable is not simple because going from
    false->true is fine but it should not be possible to remove the
    requirement to verify a signature.  But it can be useful to inspect the
    value of this parameter from userspace, so change the permissions to
    make a read-only file in sysfs.

    Signed-off-by: John Keeping <john@metanate.com>
    Signed-off-by: Mike Snitzer <snitzer@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a0753cb6276ccd5b2008d2d98f7fc05fe4ff25c
Author: Daniel Borkmann <daniel@iogearbox.net>
Date:   Tue May 4 08:58:25 2021 +0000

    bpf, selftests: Adjust few selftest result_unpriv outcomes

    [ Upstream commit 1bad6fd52be4ce12d207e2820ceb0f29ab31fc53 ]

    Given we don't need to simulate the speculative domain for registers with
    immediates anymore since the verifier uses direct imm-based rewrites instead
    of having to mask, we can also lift a few cases that were previously rejected.

    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 050a7e2a54e5f77625fd430a65133ee862c3df40
Author: Tiezhu Yang <yangtiezhu@loongson.cn>
Date:   Sat May 15 19:02:01 2021 +0800

    MIPS: Fix kernel hang under FUNCTION_GRAPH_TRACER and PREEMPT_TRACER

    [ Upstream commit 78cf0eb926cb1abeff2106bae67752e032fe5f3e ]

    When update the latest mainline kernel with the following three configs,
    the kernel hangs during startup:

    (1) CONFIG_FUNCTION_GRAPH_TRACER=y
    (2) CONFIG_PREEMPT_TRACER=y
    (3) CONFIG_FTRACE_STARTUP_TEST=y

    When update the latest mainline kernel with the above two configs (1)
    and (2), the kernel starts normally, but it still hangs when execute
    the following command:

    echo "function_graph" > /sys/kernel/debug/tracing/current_tracer

    Without CONFIG_PREEMPT_TRACER=y, the above two kinds of kernel hangs
    disappeared, so it seems that CONFIG_PREEMPT_TRACER has some influences
    with function_graph tracer at the first glance.

    I use ejtag to find out the epc address is related with preempt_enable()
    in the file arch/mips/lib/mips-atomic.c, because function tracing can
    trace the preempt_{enable,disable} calls that are traced, replace them
    with preempt_{enable,disable}_notrace to prevent function tracing from
    going into an infinite loop, and then it can fix the kernel hang issue.

    By the way, it seems that this commit is a complement and improvement of
    commit f93a1a00f2bd ("MIPS: Fix crash that occurs when function tracing
    is enabled").
```

```
    Signed-off-by: Tiezhu Yang <yangtiezhu@loongson.cn>
    Cc: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 268323c154abdaecaf20b62507964a8e1df1bd54
Author: Hannes Reinecke <hare@suse.de>
Date:   Fri May 21 10:23:46 2021 +0200

    nvme-fabrics: decode host pathing error for connect

    [ Upstream commit 4d9442bf263ac45d495bb7ecf75009e59c0622b2 ]

    Add an additional decoding for 'host pathing error' during connect.

    Signed-off-by: Hannes Reinecke <hare@suse.de>
    Reviewed-by: Sagi Grimberg <sagi@grimberg.me>
    Reviewed-by: Chaitanya Kulkarni <chaitanya.kulkarni@wdc.com>
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3dc1e208ca8dea01870ce341eafe63a924baf049
Author: Nick Desaulniers <ndesaulniers@google.com>
Date:   Thu Mar 11 17:09:41 2021 -0800

    Makefile: LTO: have linker check -Wframe-larger-than

    [ Upstream commit 24845dcb170e16b3100bd49743687648c71387ae ]

    -Wframe-larger-than= requires stack frame information, which the
    frontend cannot provide. This diagnostic is emitted late during
    compilation once stack frame size is available.

    When building with LTO, the frontend simply lowers C to LLVM IR and does
    not have stack frame information, so it cannot emit this diagnostic.
    When the linker drives LTO, it restarts optimizations and lowers LLVM IR
    to object code. At that point, it has stack frame information but
    doesn't know to check for a specific max stack frame size.

    I consider this a bug in LLVM that we need to fix. There are some
    details we're working out related to LTO such as which value to use when
    there are multiple different values specified per TU, or how to
    propagate these to compiler synthesized routines properly, if at all.

    Until it's fixed, ensure we don't miss these. At that point we can wrap
    this in a compiler version guard or revert this based on the minimum
    support version of Clang.

    The error message is not generated during link:
      LTO     vmlinux.o
    ld.lld: warning: stack size limit exceeded (8224) in foobarbaz

    Cc: Sami Tolvanen <samitolvanen@google.com>
    Reported-by: Candle Sun <candlesea@gmail.com>
    Suggested-by: Fangrui Song <maskray@google.com>
    Signed-off-by: Nick Desaulniers <ndesaulniers@google.com>
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20210312010942.1546679-1-ndesaulniers@google.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f34aff56ff0dffef95a5e21c604b77ff48de35b
Author: George McCollister <george.mccollister@gmail.com>
Date:   Mon May 24 15:29:53 2021 -0500
```

```
    net: dsa: microchip: enable phy errata workaround on 9567

    [ Upstream commit 8c42a49738f16af0061f9ae5c2f5a955f268d9e3 ]

    Also enable phy errata workaround on 9567 since has the same errata as
    the 9477 according to the manufacture's documentation.

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

commit 42e7fa0c2a4cb4834c0df0644ca08d9ea5e648bf
Author: Saubhik Mukherjee <saubhik.mukherjee@gmail.com>
Date:   Mon May 24 19:07:12 2021 +0530

    net: appletalk: cops: Fix data race in cops_probe1

    [ Upstream commit a4dd4fc6105e54393d637450a11d4cddb5fabc4f ]

    In cops_probe1(), there is a write to dev->base_addr after requesting an
    interrupt line and registering the interrupt handler cops_interrupt().
    The handler might be called in parallel to handle an interrupt.
    cops_interrupt() tries to read dev->base_addr leading to a potential
    data race. So write to dev->base_addr before calling request_irq().

    Found by Linux Driver Verification project (linuxtesting.org).

    Signed-off-by: Saubhik Mukherjee <saubhik.mukherjee@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f66d0d2f0334e5624bea35ad7a8cea75aabbbff
Author: Zong Li <zong.li@sifive.com>
Date:   Sat May 22 17:16:11 2021 +0800

    net: macb: ensure the device is available before accessing GEMGXL control registers

    [ Upstream commit 5eff1461a6dec84f04fafa9128548bad51d96147 ]

    If runtime power menagement is enabled, the gigabit ethernet PLL would
    be disabled after macb_probe(). During this period of time, the system
    would hang up if we try to access GEMGXL control registers.

    We can't put runtime_pm_get/runtime_pm_put/ there due to the issue of
    sleep inside atomic section (7fa2955ff70ce453 ("sh_eth: Fix sleeping
    function called from invalid context"). Add netif_running checking to
    ensure the device is available before accessing GEMGXL device.

    Changed in v2:
     - Use netif_running instead of its own flag

    Signed-off-by: Zong Li <zong.li@sifive.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 817a9bb19cc222b4e8eadfef67f51e4a46a85715
Author: Dmitry Bogdanov <d.bogdanov@yadro.com>
Date:   Thu Apr 15 23:35:54 2021 +0300

    scsi: target: qla2xxx: Wait for stop_phase1 at WWN removal

    [ Upstream commit 2ef7665dfd88830f15415ba007c7c9a46be7acd8 ]
```

```
    Target de-configuration panics at high CPU load because TPGT and WWPN can
    be removed on separate threads.

    TPGT removal requests a reset HBA on a separate thread and waits for reset
    complete (phase1). Due to high CPU load that HBA reset can be delayed for
    some time.

    WWPN removal does qlt_stop_phase2(). There it is believed that phase1 has
    already completed and thus tgt.tgt_ops is subsequently cleared. However,
    tgt.tgt_ops is needed to process incoming traffic and therefore this will
    cause one of the following panics:

    NIP qlt_reset+0x7c/0x220 [qla2xxx]
    LR  qlt_reset+0x68/0x220 [qla2xxx]
    Call Trace:
    0xc000003ffff63a78 (unreliable)
    qlt_handle_imm_notify+0x800/0x10c0 [qla2xxx]
    qlt_24xx_atio_pkt+0x208/0x590 [qla2xxx]
    qlt_24xx_process_atio_queue+0x33c/0x7a0 [qla2xxx]
    qla83xx_msix_atio_q+0x54/0x90 [qla2xxx]

    or

    NIP qlt_24xx_handle_abts+0xd0/0x2a0 [qla2xxx]
    LR  qlt_24xx_handle_abts+0xb4/0x2a0 [qla2xxx]
    Call Trace:
    qlt_24xx_handle_abts+0x90/0x2a0 [qla2xxx] (unreliable)
    qlt_24xx_process_atio_queue+0x500/0x7a0 [qla2xxx]
    qla83xx_msix_atio_q+0x54/0x90 [qla2xxx]

    or

    NIP qlt_create_sess+0x90/0x4e0 [qla2xxx]
    LR  qla24xx_do_nack_work+0xa8/0x180 [qla2xxx]
    Call Trace:
    0xc0000000348fba30 (unreliable)
    qla24xx_do_nack_work+0xa8/0x180 [qla2xxx]
    qla2x00_do_work+0x674/0xbf0 [qla2xxx]
    qla2x00_iocb_work_fn

    The patch fixes the issue by serializing qlt_stop_phase1() and
    qlt_stop_phase2() functions to make WWPN removal wait for phase1
    completion.

    Link: https://lore.kernel.org/r/20210415203554.27890-1-d.bogdanov@yadro.com
    Reviewed-by: Roman Bolshakov <r.bolshakov@yadro.com>
    Signed-off-by: Dmitry Bogdanov <d.bogdanov@yadro.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b392e2fb0ee5c30c4f7c4f171474f617a2529f1d
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed May 19 21:05:19 2021 +0800

    scsi: hisi_sas: Drop free_irq() of devm_request_irq() allocated irq

    [ Upstream commit 7907a021e4bbfa29cccacd2ba2dade894d9a7d4c ]

    irqs allocated with devm_request_irq() should not be freed using
    free_irq(). Doing so causes a dangling pointer and a subsequent double
    free.

    Link: https://lore.kernel.org/r/20210519130519.2661938-1-yangyingliang@huawei.com
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Acked-by: John Garry <john.garry@huawei.com>
```

```
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 87bf5d1a72ca03f4f4d0c6112d3055d6d72f551d
Author: Matt Wang <wwentao@vmware.com>
Date:   Wed May 19 09:49:32 2021 +0000

    scsi: vmw_pvscsi: Set correct residual data length

    [ Upstream commit e662502b3a782d479e67736a5a1c169a703d853a ]

    Some commands (such as INQUIRY) may return less data than the initiator
    requested. To avoid conducting useless information, set the right residual
    count to make upper layer aware of this.

    Before (INQUIRY PAGE 0xB0 with 128B buffer):

    $ sg_raw -r 128 /dev/sda 12 01 B0 00 80 00
    SCSI Status: Good

    Received 128 bytes of data:
     00 00 b0 00 3c 01 00 00 00 00 00 00 00 00 00 00 ...<............
     10 00 00 00 00 00 01 00 00 00 00 00 40 00 00 08 ............@....
     20 80 00 00 00 00 00 00 00 00 00 20 00 00 00 00 ..........  .....
     30 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
     40 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
     50 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
     60 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
     70 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................

    After:

    $ sg_raw -r 128 /dev/sda 12 01 B0 00 80 00
    SCSI Status: Good

    Received 64 bytes of data:
    00 00 b0 00 3c 01 00 00 00 00 00 00 00 00 00 00 ...<............
    10 00 00 00 00 00 01 00 00 00 00 00 40 00 00 08 ............@....
    20 80 00 00 00 00 00 00 00 00 00 20 00 00 00 00 ..........  .....
    30 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................

    [mkp: clarified description]

    Link: https://lore.kernel.org/r/03C41093-B62E-43A2-913E-CFC92F1C70C3@vmware.com
    Signed-off-by: Matt Wang <wwentao@vmware.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e573390098969ee6889686bcf91e11cbc2d0558a
Author: Javed Hasan <jhasan@marvell.com>
Date:   Tue May 18 23:14:16 2021 -0700

    scsi: bnx2fc: Return failure if io_req is already in ABTS processing

    [ Upstream commit 122c81c563b0c1c6b15ff76a9159af5ee1f21563 ]

    Return failure from bnx2fc_eh_abort() if io_req is already in ABTS
    processing.

    Link: https://lore.kernel.org/r/20210519061416.19321-1-jhasan@marvell.com
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Javed Hasan <jhasan@marvell.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 81c4d1d83f88e15b26f4522a35cba6ffd8c5dfdd
Author: Íñigo Huguet <ihuguet@redhat.com>
Date:   Fri May 21 16:38:35 2021 +0200

    net:sfc: fix non-freed irq in legacy irq mode

    [ Upstream commit 8f03eeb6e0a0a0b8d617ee0a4bce729e47130036 ]

    SFC driver can be configured via modparam to work using MSI-X, MSI or
    legacy IRQ interrupts. In the last one, the interrupt was not properly
    released on module remove.

    It was not freed because the flag irqs_hooked was not set during
    initialization in the case of using legacy IRQ.

    Example of (trimmed) trace during module remove without this fix:

    remove_proc_entry: removing non-empty directory 'irq/125', leaking at least
'0000:3b:00.1'
    WARNING: CPU: 39 PID: 3658 at fs/proc/generic.c:715 remove_proc_entry+0x15c/0x170
    ...trimmed...
    Call Trace:
     unregister_irq_proc+0xe3/0x100
     free_desc+0x29/0x70
     irq_free_descs+0x47/0x70
     mp_unmap_irq+0x58/0x60
     acpi_unregister_gsi_ioapic+0x2a/0x40
     acpi_pci_irq_disable+0x78/0xb0
     pci_disable_device+0xd1/0x100
     efx_pci_remove+0xa1/0x1e0 [sfc]
     pci_device_remove+0x38/0xa0
     __device_release_driver+0x177/0x230
     driver_detach+0xcb/0x110
     bus_remove_driver+0x58/0xd0
     pci_unregister_driver+0x2a/0xb0
     efx_exit_module+0x24/0xf40 [sfc]
     __do_sys_delete_module.constprop.0+0x171/0x280
     ? exit_to_user_mode_prepare+0x83/0x1d0
     do_syscall_64+0x3d/0x80
     entry_SYSCALL_64_after_hwframe+0x44/0xae
    RIP: 0033:0x7f9f9385800b
    ...trimmed...

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

commit 495d0518640a8a21dc8d8686af56098028f3e6e5
Author: Rao Shoaib <rao.shoaib@oracle.com>
Date:   Fri May 21 11:08:06 2021 -0700

    RDS tcp loopback connection can hang

    [ Upstream commit aced3ce57cd37b5ca332bcacd370d01f5a8c5371 ]

    When TCP is used as transport and a program on the
    system connects to RDS port 16385, connection is
    accepted but denied per the rules of RDS. However,
    RDS connections object is left in the list. Next
    loopback connection will select that connection
    object as it is at the head of list. The connection
    attempt will hang as the connection object is set
    to connect over TCP which is not allowed
```

```
    The issue can be reproduced easily, use rds-ping
    to ping a local IP address. After that use any
    program like ncat to connect to the same IP
    address and port 16385. This will hang so ctrl-c out.
    Now try rds-ping, it will hang.

    To fix the issue this patch adds checks to disallow
    the connection object creation and destroys the
    connection object.

    Signed-off-by: Rao Shoaib <rao.shoaib@oracle.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ed296c47e5c545a26bb774ec7a87e5369e6ddf9
Author: Zheyu Ma <zheyuma97@gmail.com>
Date:   Thu May 20 12:32:36 2021 +0000

    net/qla3xxx: fix schedule while atomic in ql_sem_spinlock

    [ Upstream commit 13a6f3153922391e90036ba2267d34eed63196fc ]

    When calling the 'ql_sem_spinlock', the driver has already acquired the
    spin lock, so the driver should not call 'ssleep' in atomic context.

    This bug can be fixed by using 'mdelay' instead of 'ssleep'.

    The KASAN's log reveals it:

    [    3.238124 ] BUG: scheduling while atomic: swapper/0/1/0x00000002
    [    3.238748 ] 2 locks held by swapper/0/1:
    [    3.239151 ]  #0: ffff88810177b240 (&dev->mutex){....}-{3:3}, at:
    __device_driver_lock+0x41/0x60
    [    3.240026 ]  #1: ffff888107c60e28 (&qdev->hw_lock){....}-{2:2}, at:
    ql3xxx_probe+0x2aa/0xea0
    [    3.240873 ] Modules linked in:
    [    3.241187 ] irq event stamp: 460854
    [    3.241541 ] hardirqs last  enabled at (460853): [<ffffffff843051bf>]
    _raw_spin_unlock_irqrestore+0x4f/0x70
    [    3.242245 ] hardirqs last disabled at (460854): [<ffffffff843058ca>]
    _raw_spin_lock_irqsave+0x2a/0x70
    [    3.242245 ] softirqs last  enabled at (446076): [<ffffffff846002e4>]
    __do_softirq+0x2e4/0x4b1
    [    3.242245 ] softirqs last disabled at (446069): [<ffffffff811ba5e0>]
    irq_exit_rcu+0x100/0x110
    [    3.242245 ] Preemption disabled at:
    [    3.242245 ] [<ffffffff828ca5ba>] ql3xxx_probe+0x2aa/0xea0
    [    3.242245 ] Kernel panic - not syncing: scheduling while atomic
    [    3.242245 ] CPU: 2 PID: 1 Comm: swapper/0 Not tainted
    5.13.0-rc1-00145
    -gee7dc339169-dirty #16
    [    3.242245 ] Call Trace:
    [    3.242245 ]  dump_stack+0xba/0xf5
    [    3.242245 ]  ? ql3xxx_probe+0x1f0/0xea0
    [    3.242245 ]  panic+0x15a/0x3f2
    [    3.242245 ]  ? vprintk+0x76/0x150
    [    3.242245 ]  ? ql3xxx_probe+0x2aa/0xea0
    [    3.242245 ]  __schedule_bug+0xae/0xe0
    [    3.242245 ]  __schedule+0x72e/0xa00
    [    3.242245 ]  schedule+0x43/0xf0
    [    3.242245 ]  schedule_timeout+0x28b/0x500
    [    3.242245 ]  ? del_timer_sync+0xf0/0xf0
    [    3.242245 ]  ? msleep+0x2f/0x70
    [    3.242245 ]  msleep+0x59/0x70
    [    3.242245 ]  ql3xxx_probe+0x307/0xea0
```

```
[    3.242245 ]  ? _raw_spin_unlock_irqrestore+0x3a/0x70
[    3.242245 ]  ? pci_device_remove+0x110/0x110
[    3.242245 ]  local_pci_probe+0x45/0xa0
[    3.242245 ]  pci_device_probe+0x12b/0x1d0
[    3.242245 ]  really_probe+0x2a9/0x610
[    3.242245 ]  driver_probe_device+0x90/0x1d0
[    3.242245 ]  ? mutex_lock_nested+0x1b/0x20
[    3.242245 ]  device_driver_attach+0x68/0x70
[    3.242245 ]  __driver_attach+0x124/0x1b0
[    3.242245 ]  ? device_driver_attach+0x70/0x70
[    3.242245 ]  bus_for_each_dev+0xbb/0x110
[    3.242245 ]  ? rdinit_setup+0x45/0x45
[    3.242245 ]  driver_attach+0x27/0x30
[    3.242245 ]  bus_add_driver+0x1eb/0x2a0
[    3.242245 ]  driver_register+0xa9/0x180
[    3.242245 ]  __pci_register_driver+0x82/0x90
[    3.242245 ]  ? yellowfin_init+0x25/0x25
[    3.242245 ]  ql3xxx_driver_init+0x23/0x25
[    3.242245 ]  do_one_initcall+0x7f/0x3d0
[    3.242245 ]  ? rdinit_setup+0x45/0x45
[    3.242245 ]  ? rcu_read_lock_sched_held+0x4f/0x80
[    3.242245 ]  kernel_init_freeable+0x2aa/0x301
[    3.242245 ]  ? rest_init+0x2c0/0x2c0
[    3.242245 ]  kernel_init+0x18/0x190
[    3.242245 ]  ? rest_init+0x2c0/0x2c0
[    3.242245 ]  ? rest_init+0x2c0/0x2c0
[    3.242245 ]  ret_from_fork+0x1f/0x30
[    3.242245 ] Dumping ftrace buffer:
[    3.242245 ]    (ftrace buffer empty)
[    3.242245 ] Kernel Offset: disabled
[    3.242245 ] Rebooting in 1 seconds.
```

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

commit 8b30908df8d53db5ccb7f1bcd1c97d3aa75bbd62
Author: Sergey Senozhatsky <senozhatsky@chromium.org>
Date: Thu May 20 19:14:22 2021 +0900

    wq: handle VM suspension in stall detection

    [Upstream commit 940d71c6462e8151c78f28e4919aa8882ff2054e]

    If VCPU is suspended (VM suspend) in wq_watchdog_timer_fn() then
    once this VCPU resumes it will see the new jiffies value, while it
    may take a while before IRQ detects PVCLOCK_GUEST_STOPPED on this
    VCPU and updates all the watchdogs via pvclock_touch_watchdogs().
    There is a small chance of misreported WQ stalls in the meantime,
    because new jiffies is time_after() old 'ts + thresh'.

    wq_watchdog_timer_fn()
    {
            for_each_pool(pool, pi) {
                    if (time_after(jiffies, ts + thresh)) {
                            pr_emerg("BUG: workqueue lockup - pool");
                    }
            }
    }

    Save jiffies at the beginning of this function and use that value
    for stall detection. If VM gets suspended then we continue using
    "old" jiffies value and old WQ touch timestamps. If IRQ at some
    point restarts the stall detection cycle (pvclock_touch_watchdogs())

```
    then old jiffies will always be before new 'ts + thresh'.

    Signed-off-by: Sergey Senozhatsky <senozhatsky@chromium.org>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac91f6bdc0c5863834a2fcffab706f821aec651e
Author: Shakeel Butt <shakeelb@google.com>
Date:   Wed May 12 13:19:46 2021 -0700

    cgroup: disable controllers at parse time

    [ Upstream commit 45e1ba40837ac2f6f4d4716bddb8d44bd7e4a251 ]

    This patch effectively reverts the commit a3e72739b7a7 ("cgroup: fix
    too early usage of static_branch_disable()"). The commit 6041186a3258
    ("init: initialize jump labels before command line option parsing") has
    moved the jump_label_init() before parse_args() which has made the
    commit a3e72739b7a7 unnecessary. On the other hand there are
    consequences of disabling the controllers later as there are subsystems
    doing the controller checks for different decisions. One such incident
    is reported [1] regarding the memory controller and its impact on memory
    reclaim code.

    [1] https://lore.kernel.org/linux-mm/921e53f3-4b13-aab8-4a9e-e83ff15371e4@nec.com

    Signed-off-by: Shakeel Butt <shakeelb@google.com>
    Reported-by: NOMURA JUNICHI(野村 淳一) <junichi.nomura@nec.com>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Tested-by: Jun'ichi Nomura <junichi.nomura@nec.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a244250207f5e4bacd38533a306e8a8d49f48685
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Mon May 17 12:04:13 2021 +0300

    net: mdiobus: get rid of a BUG_ON()

    [ Upstream commit 1dde47a66d4fb181830d6fa000e5ea86907b639e ]

    We spotted a bug recently during a review where a driver was
    unregistering a bus that wasn't registered, which would trigger this
    BUG_ON().  Let's handle that situation more gracefully, and just print
    a warning and return.

    Reported-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Reviewed-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76cc8e04f38c2bbfcba07f62864a011f142bd40c
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Mon May 17 16:38:09 2021 +0200

    netlink: disable IRQs for netlink_lock_table()

    [ Upstream commit 1d482e666b8e74c7555dbdfbfb77205eeed3ff2d ]

    Syzbot reports that in mac80211 we have a potential deadlock
    between our "local->stop_queue_reasons_lock" (spinlock) and
    netlink's nl_table_lock (rwlock). This is because there's at
    least one situation in which we might try to send a netlink
    message with this spinlock held while it is also possible to
```

```
    take the spinlock from a hardirq context, resulting in the
    following deadlock scenario reported by lockdep:

          CPU0                    CPU1
          ----                    ----
      lock(nl_table_lock);
                                  local_irq_disable();
                                  lock(&local->queue_stop_reason_lock);
                                  lock(nl_table_lock);
      <Interrupt>
        lock(&local->queue_stop_reason_lock);

    This seems valid, we can take the queue_stop_reason_lock in
    any kind of context ("CPU0"), and call ieee80211_report_ack_skb()
    with the spinlock held and IRQs disabled ("CPU1") in some
    code path (ieee80211_do_stop() via ieee80211_free_txskb()).

    Short of disallowing netlink use in scenarios like these
    (which would be rather complex in mac80211's case due to
    the deep callchain), it seems the only fix for this is to
    disable IRQs while nl_table_lock is held to avoid hitting
    this scenario, this disallows the "CPU0" portion of the
    reported deadlock.

    Note that the writer side (netlink_table_grab()) already
    disables IRQs for this lock.

    Unfortunately though, this seems like a huge hammer, and
    maybe the whole netlink table locking should be reworked.

    Reported-by: syzbot+69ff9dff50dcfe14ddd4@syzkaller.appspotmail.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf3f90b7b5abc2154c3c611b6846b265932f47ce
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Mon May 17 16:13:35 2021 +0200

    bonding: init notify_work earlier to avoid uninitialized use

    [ Upstream commit 35d96e631860226d5dc4de0fad0a415362ec2457 ]

    If bond_kobj_init() or later kzalloc() in bond_alloc_slave() fail,
    then we call kobject_put() on the slave->kobj. This in turn calls
    the release function slave_kobj_release() which will always try to
    cancel_delayed_work_sync(&slave->notify_work), which shouldn't be
    done on an uninitialized work struct.

    Always initialize the work struct earlier to avoid problems here.

    Syzbot bisected this down to a completely pointless commit, some
    fault injection may have been at work here that caused the alloc
    failure in the first place, which may interact badly with bisect.

    Reported-by: syzbot+bfda097c12a00c8cae67@syzkaller.appspotmail.com
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Acked-by: Jay Vosburgh <jay.vosburgh@canonical.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6249193e03709ea625e10706ecaf17fea0427d3d
Author: Zheyu Ma <zheyuma97@gmail.com>
Date:   Sun May 16 07:11:40 2021 +0000
```

```
    isdn: mISDN: netjet: Fix crash in nj_probe:

    [ Upstream commit 9f6f852550d0e1b7735651228116ae9d300f69b3 ]

    'nj_setup' in netjet.c might fail with -EIO and in this case
    'card->irq' is initialized and is bigger than zero. A subsequent call to
    'nj_release' will free the irq that has not been requested.

    Fix this bug by deleting the previous assignment to 'card->irq' and just
    keep the assignment before 'request_irq'.

    The KASAN's log reveals it:

    [    3.354615 ] WARNING: CPU: 0 PID: 1 at kernel/irq/manage.c:1826
    free_irq+0x100/0x480
    [    3.355112 ] Modules linked in:
    [    3.355310 ] CPU: 0 PID: 1 Comm: swapper/0 Not tainted
    5.13.0-rc1-00144-g25a1298726e #13
    [    3.355816 ] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS
    rel-1.12.0-59-gc9ba5276e321-prebuilt.qemu.org 04/01/2014
    [    3.356552 ] RIP: 0010:free_irq+0x100/0x480
    [    3.356820 ] Code: 6e 08 74 6f 4d 89 f4 e8 5e ac 09 00 4d 8b 74 24 18
    4d 85 f6 75 e3 e8 4f ac 09 00 8b 75 c8 48 c7 c7 78 c1 2e 85 e8 e0 cf f5
    ff <0f> 0b 48 8b 75 c0 4c 89 ff e8 72 33 0b 03 48 8b 43 40 4c 8b a0 80
    [    3.358012 ] RSP: 0000:ffffc90000017b48 EFLAGS: 00010082
    [    3.358357 ] RAX: 0000000000000000 RBX: ffff888104dc8000 RCX:
    0000000000000000
    [    3.358814 ] RDX: ffff8881003c8000 RSI: ffffffff8124a9e6 RDI:
    00000000ffffffff
    [    3.359272 ] RBP: ffffc90000017b88 R08: 0000000000000000 R09:
    0000000000000000
    [    3.359732 ] R10: ffffc900000179f0 R11: 0000000000001d04 R12:
    0000000000000000
    [    3.360195 ] R13: ffff888107dc6000 R14: ffff888107dc6928 R15:
    ffff888104dc80a8
    [    3.360652 ] FS:  0000000000000000(0000) GS:ffff88817bc00000(0000)
    knlGS:0000000000000000
    [    3.361170 ] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [    3.361538 ] CR2: 0000000000000000 CR3: 000000000582e000 CR4:
    00000000000006f0
    [    3.362003 ] DR0: 0000000000000000 DR1: 0000000000000000 DR2:
    0000000000000000
    [    3.362175 ] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7:
    0000000000000400
    [    3.362175 ] Call Trace:
    [    3.362175 ]  nj_release+0x51/0x1e0
    [    3.362175 ]  nj_probe+0x450/0x950
    [    3.362175 ]  ? pci_device_remove+0x110/0x110
    [    3.362175 ]  local_pci_probe+0x45/0xa0
    [    3.362175 ]  pci_device_probe+0x12b/0x1d0
    [    3.362175 ]  really_probe+0x2a9/0x610
    [    3.362175 ]  driver_probe_device+0x90/0x1d0
    [    3.362175 ]  ? mutex_lock_nested+0x1b/0x20
    [    3.362175 ]  device_driver_attach+0x68/0x70
    [    3.362175 ]  __driver_attach+0x124/0x1b0
    [    3.362175 ]  ? device_driver_attach+0x70/0x70
    [    3.362175 ]  bus_for_each_dev+0xbb/0x110
    [    3.362175 ]  ? rdinit_setup+0x45/0x45
    [    3.362175 ]  driver_attach+0x27/0x30
    [    3.362175 ]  bus_add_driver+0x1eb/0x2a0
    [    3.362175 ]  driver_register+0xa9/0x180
    [    3.362175 ]  __pci_register_driver+0x82/0x90
    [    3.362175 ]  ? w6692_init+0x38/0x38
    [    3.362175 ]  nj_init+0x36/0x38
    [    3.362175 ]  do_one_initcall+0x7f/0x3d0
```

```
[    3.362175 ]  ? rdinit_setup+0x45/0x45
[    3.362175 ]  ? rcu_read_lock_sched_held+0x4f/0x80
[    3.362175 ]  kernel_init_freeable+0x2aa/0x301
[    3.362175 ]  ? rest_init+0x2c0/0x2c0
[    3.362175 ]  kernel_init+0x18/0x190
[    3.362175 ]  ? rest_init+0x2c0/0x2c0
[    3.362175 ]  ? rest_init+0x2c0/0x2c0
[    3.362175 ]  ret_from_fork+0x1f/0x30
[    3.362175 ] Kernel panic - not syncing: panic_on_warn set ...
[    3.362175 ] CPU: 0 PID: 1 Comm: swapper/0 Not tainted
5.13.0-rc1-00144-g25a1298726e #13
[    3.362175 ] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS
rel-1.12.0-59-gc9ba5276e321-prebuilt.qemu.org 04/01/2014
[    3.362175 ] Call Trace:
[    3.362175 ]  dump_stack+0xba/0xf5
[    3.362175 ]  ? free_irq+0x100/0x480
[    3.362175 ]  panic+0x15a/0x3f2
[    3.362175 ]  ? __warn+0xf2/0x150
[    3.362175 ]  ? free_irq+0x100/0x480
[    3.362175 ]  __warn+0x108/0x150
[    3.362175 ]  ? free_irq+0x100/0x480
[    3.362175 ]  report_bug+0x119/0x1c0
[    3.362175 ]  handle_bug+0x3b/0x80
[    3.362175 ]  exc_invalid_op+0x18/0x70
[    3.362175 ]  asm_exc_invalid_op+0x12/0x20
[    3.362175 ] RIP: 0010:free_irq+0x100/0x480
[    3.362175 ] Code: 6e 08 74 6f 4d 89 f4 e8 5e ac 09 00 4d 8b 74 24 18
4d 85 f6 75 e3 e8 4f ac 09 00 8b 75 c8 48 c7 c7 78 c1 2e 85 e8 e0 cf f5
ff <0f> 0b 48 8b 75 c0 4c 89 ff e8 72 33 0b 03 48 8b 43 40 4c 8b a0 80
[    3.362175 ] RSP: 0000:ffffc90000017b48 EFLAGS: 00010082
[    3.362175 ] RAX: 0000000000000000 RBX: ffff888104dc8000 RCX:
0000000000000000
[    3.362175 ] RDX: ffff8881003c8000 RSI: ffffffff8124a9e6 RDI:
00000000ffffffff
[    3.362175 ] RBP: ffffc90000017b88 R08: 0000000000000000 R09:
0000000000000000
[    3.362175 ] R10: ffffc900000179f0 R11: 0000000000001d04 R12:
0000000000000000
[    3.362175 ] R13: ffff888107dc6000 R14: ffff888107dc6928 R15:
ffff888104dc80a8
[    3.362175 ]  ? vprintk+0x76/0x150
[    3.362175 ]  ? free_irq+0x100/0x480
[    3.362175 ]  nj_release+0x51/0x1e0
[    3.362175 ]  nj_probe+0x450/0x950
[    3.362175 ]  ? pci_device_remove+0x110/0x110
[    3.362175 ]  local_pci_probe+0x45/0xa0
[    3.362175 ]  pci_device_probe+0x12b/0x1d0
[    3.362175 ]  really_probe+0x2a9/0x610
[    3.362175 ]  driver_probe_device+0x90/0x1d0
[    3.362175 ]  ? mutex_lock_nested+0x1b/0x20
[    3.362175 ]  device_driver_attach+0x68/0x70
[    3.362175 ]  __driver_attach+0x124/0x1b0
[    3.362175 ]  ? device_driver_attach+0x70/0x70
[    3.362175 ]  bus_for_each_dev+0xbb/0x110
[    3.362175 ]  ? rdinit_setup+0x45/0x45
[    3.362175 ]  driver_attach+0x27/0x30
[    3.362175 ]  bus_add_driver+0x1eb/0x2a0
[    3.362175 ]  driver_register+0xa9/0x180
[    3.362175 ]  __pci_register_driver+0x82/0x90
[    3.362175 ]  ? w6692_init+0x38/0x38
[    3.362175 ]  nj_init+0x36/0x38
[    3.362175 ]  do_one_initcall+0x7f/0x3d0
[    3.362175 ]  ? rdinit_setup+0x45/0x45
[    3.362175 ]  ? rcu_read_lock_sched_held+0x4f/0x80
[    3.362175 ]  kernel_init_freeable+0x2aa/0x301
```

```
     [    3.362175 ]  ? rest_init+0x2c0/0x2c0
     [    3.362175 ]  kernel_init+0x18/0x190
     [    3.362175 ]  ? rest_init+0x2c0/0x2c0
     [    3.362175 ]  ? rest_init+0x2c0/0x2c0
     [    3.362175 ]  ret_from_fork+0x1f/0x30
     [    3.362175 ] Dumping ftrace buffer:
     [    3.362175 ]    (ftrace buffer empty)
     [    3.362175 ] Kernel Offset: disabled
     [    3.362175 ] Rebooting in 1 seconds..

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

commit 0a42b3feff92d80b561e0d78c1d327d61c91a168
Author: Li Jun <jun.li@nxp.com>
Date:   Fri May 14 18:59:44 2021 +0800

    usb: chipidea: udc: assign interrupt number to USB gadget structure

    [ Upstream commit 9e3927f6373da54cb17e17f4bd700907e1123d2f ]

    Chipidea also need sync interrupt before unbind the udc while
    gadget remove driver, otherwise setup irq handling may happen
    while unbind, see below dump generated from android function
    switch stress test:

    [ 4703.503056] android_work: sent uevent USB_STATE=CONNECTED
    [ 4703.514642] android_work: sent uevent USB_STATE=DISCONNECTED
    [ 4703.651339] android_work: sent uevent USB_STATE=CONNECTED
    [ 4703.661806] init: Control message: Processed ctl.stop for 'adbd' from pid: 561
(system_server)
    [ 4703.673469] init: processing action (init.svc.adbd=stopped) from
(/system/etc/init/hw/init.usb.configfs.rc:14)
    [ 4703.676451] Unable to handle kernel read from unreadable memory at virtual address
0000000000000090
    [ 4703.676454] Mem abort info:
    [ 4703.676458]   ESR = 0x96000004
    [ 4703.676461]   EC = 0x25: DABT (current EL), IL = 32 bits
    [ 4703.676464]   SET = 0, FnV = 0
    [ 4703.676466]   EA = 0, S1PTW = 0
    [ 4703.676468] Data abort info:
    [ 4703.676471]   ISV = 0, ISS = 0x00000004
    [ 4703.676473]   CM = 0, WnR = 0
    [ 4703.676478] user pgtable: 4k pages, 48-bit VAs, pgdp=000000004a867000
    [ 4703.676481] [0000000000000090] pgd=0000000000000000, p4d=0000000000000000
    [ 4703.676503] Internal error: Oops: 96000004 [#1] PREEMPT SMP
    [ 4703.758297] Modules linked in: synaptics_dsx_i2c moal(O) mlan(O)
    [ 4703.764327] CPU: 0 PID: 235 Comm: lmkd Tainted: G        W  O      5.10.9-00001-
g3f5fd8487c38-dirty #63
    [ 4703.773720] Hardware name: NXP i.MX8MNano EVK board (DT)
    [ 4703.779033] pstate: 60400085 (nZCv daIf +PAN -UAO -TCO BTYPE=--)
    [ 4703.785046] pc : _raw_write_unlock_bh+0xc0/0x2c8
    [ 4703.789667] lr : android_setup+0x4c/0x168
    [ 4703.793676] sp : ffff80001256bd80
    [ 4703.796989] x29: ffff80001256bd80 x28: 00000000000000a8
    [ 4703.802304] x27: ffff800012470000 x26: ffff80006d923000
    [ 4703.807616] x25: ffff800012471000 x24: ffff00000b091140
    [ 4703.812929] x23: ffff0000077dbd38 x22: ffff0000077da490
    [ 4703.818242] x21: ffff80001256be30 x20: 0000000000000000
    [ 4703.823554] x19: 0000000000000080 x18: ffff800012561048
    [ 4703.828867] x17: 0000000000000000 x16: 0000000000000039
    [ 4703.834180] x15: ffff8000106ad258 x14: ffff80001194c277
    [ 4703.839493] x13: 0000000000003934 x12: 0000000000000000
```

```
     [ 4703.844805] x11: 0000000000000000 x10: 0000000000000001
     [ 4703.850117] x9 : 0000000000000000 x8 : 0000000000000090
     [ 4703.855429] x7 : 6f72646e61203a70 x6 : ffff8000124f2450
     [ 4703.860742] x5 : ffffffffffffffff x4 : 0000000000000009
     [ 4703.866054] x3 : ffff8000108a290c x2 : ffff00007fb3a9c8
     [ 4703.871367] x1 : 0000000000000000 x0 : 0000000000000090
     [ 4703.876681] Call trace:
     [ 4703.879129]  _raw_write_unlock_bh+0xc0/0x2c8
     [ 4703.883397]  android_setup+0x4c/0x168
     [ 4703.887059]  udc_irq+0x824/0xa9c
     [ 4703.890287]  ci_irq+0x124/0x148
     [ 4703.893429]  __handle_irq_event_percpu+0x84/0x268
     [ 4703.898131]  handle_irq_event+0x64/0x14c
     [ 4703.902054]  handle_fasteoi_irq+0x110/0x210
     [ 4703.906236]  __handle_domain_irq+0x8c/0xd4
     [ 4703.910332]  gic_handle_irq+0x6c/0x124
     [ 4703.914081]  el1_irq+0xdc/0x1c0
     [ 4703.917221]  _raw_spin_unlock_irq+0x20/0x54
     [ 4703.921405]  finish_task_switch+0x84/0x224
     [ 4703.925502]  __schedule+0x4a4/0x734
     [ 4703.928990]  schedule+0xa0/0xe8
     [ 4703.932132]  do_notify_resume+0x150/0x184
     [ 4703.936140]  work_pending+0xc/0x40c
     [ 4703.939633] Code: d5384613 521b0a69 d5184609 f9800111 (885ffd01)
     [ 4703.945732] ---[ end trace ba5c1875ae49d53c ]---
     [ 4703.950350] Kernel panic - not syncing: Oops: Fatal exception in interrupt
     [ 4703.957223] SMP: stopping secondary CPUs
     [ 4703.961151] Kernel Offset: disabled
     [ 4703.964638] CPU features: 0x0240002,2000200c
     [ 4703.968905] Memory Limit: none
     [ 4703.971963] Rebooting in 5 seconds..

    Tested-by: faqiang.zhu <faqiang.zhu@nxp.com>
    Signed-off-by: Li Jun <jun.li@nxp.com>
    Link: https://lore.kernel.org/r/1620989984-7653-1-git-send-email-jun.li@nxp.com
    Signed-off-by: Peter Chen <peter.chen@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cec64306c9873a64c5cf1d49602bb4ff5e8b55d2
Author: Chunyan Zhang <chunyan.zhang@unisoc.com>
Date:   Wed May 12 17:35:34 2021 +0800

    spi: sprd: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 7907cad7d07e0055789ec0c534452f19dfe1fc80 ]

    MODULE_DEVICE_TABLE is used to extract the device information out of the
    driver and builds a table when being compiled. If using this macro,
    kernel can find the driver if available when the device is plugged in,
    and then loads that driver and initializes the device.

    Signed-off-by: Chunyan Zhang <chunyan.zhang@unisoc.com>
    Link: https://lore.kernel.org/r/20210512093534.243040-1-zhang.lyra@gmail.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 51bd6bf4719720bfff8e23fc805eaae21e34cf94
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 11:12:25 2021 +0800

    ASoC: sti-sas: add missing MODULE_DEVICE_TABLE

    [ Upstream commit e072b2671606c77538d6a4dd5dda80b508cb4816 ]

    This patch adds missing MODULE_DEVICE_TABLE definition which generates
```

```
    correct modalias for automatic loading of this driver when it is built
    as an external module.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Link: https://lore.kernel.org/r/1620789145-14936-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 320b3e0b16ddf98c63ccc74c1bfc0193cfcf34cd
Author: Eric Farman <farman@linux.ibm.com>
Date:   Tue May 11 21:56:31 2021 +0200

    vfio-ccw: Serialize FSM IDLE state with I/O completion

    [ Upstream commit 2af7a834a435460d546f0cf0a8b8e4d259f1d910 ]

    Today, the stacked call to vfio_ccw_sch_io_todo() does three things:

      1) Update a solicited IRB with CP information, and release the CP
         if the interrupt was the end of a START operation.
      2) Copy the IRB data into the io_region, under the protection of
         the io_mutex
      3) Reset the vfio-ccw FSM state to IDLE to acknowledge that
         vfio-ccw can accept more work.

    The trouble is that step 3 is (A) invoked for both solicited and
    unsolicited interrupts, and (B) sitting after the mutex for step 2.
    This second piece becomes a problem if it processes an interrupt
    for a CLEAR SUBCHANNEL while another thread initiates a START,
    thus allowing the CP and FSM states to get out of sync. That is:

        CPU 1                           CPU 2
        fsm_do_clear()
        fsm_irq()
                                        fsm_io_request()
        vfio_ccw_sch_io_todo()
                                        fsm_io_helper()

    Since the FSM state and CP should be kept in sync, let's make a
    note when the CP is released, and rely on that as an indication
    that the FSM should also be reset at the end of this routine and
    open up the device for more work.

    Signed-off-by: Eric Farman <farman@linux.ibm.com>
    Acked-by: Matthew Rosato <mjrosato@linux.ibm.com>
    Reviewed-by: Cornelia Huck <cohuck@redhat.com>
    Message-Id: <20210511195631.3995081-4-farman@linux.ibm.com>
    Signed-off-by: Cornelia Huck <cohuck@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50360f4e084ce1ce67fa5c062b418e4bcfa46223
Author: Eric Farman <farman@linux.ibm.com>
Date:   Tue May 11 21:56:30 2021 +0200

    vfio-ccw: Reset FSM state to IDLE inside FSM

    [ Upstream commit 6c02ac4c9211edabe17bda437ac97e578756f31b ]

    When an I/O request is made, the fsm_io_request() routine
    moves the FSM state from IDLE to CP_PROCESSING, and then
    fsm_io_helper() moves it to CP_PENDING if the START SUBCHANNEL
    received a cc0. Yet, the error case to go from CP_PROCESSING
    back to IDLE is done after the FSM call returns.
```

```
    Let's move this up into the FSM proper, to provide some
    better symmetry when unwinding in this case.

    Signed-off-by: Eric Farman <farman@linux.ibm.com>
    Reviewed-by: Cornelia Huck <cohuck@redhat.com>
    Acked-by: Matthew Rosato <mjrosato@linux.ibm.com>
    Message-Id: <20210511195631.3995081-3-farman@linux.ibm.com>
    Signed-off-by: Cornelia Huck <cohuck@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f5e770c0c60ab8812574a2e0d163b0efa816a825
Author: Jiri Olsa <jolsa@kernel.org>
Date:   Thu Apr 29 13:47:12 2021 +0200

    bpf: Add deny list of btf ids check for tracing programs

    [ Upstream commit 35e3815fa8102fab4dee75f3547472c66581125d ]

    The recursion check in __bpf_prog_enter and __bpf_prog_exit
    leaves some (not inlined) functions unprotected:

    In __bpf_prog_enter:
      - migrate_disable is called before prog->active is checked

    In __bpf_prog_exit:
      - migrate_enable,rcu_read_unlock_strict are called after
        prog->active is decreased

    When attaching trampoline to them we get panic like:

      traps: PANIC: double fault, error_code: 0x0
      double fault: 0000 [#1] SMP PTI
      RIP: 0010:__bpf_prog_enter+0x4/0x50
      ...
      Call Trace:
       <IRQ>
       bpf_trampoline_6442466513_0+0x18/0x1000
       migrate_disable+0x5/0x50
       __bpf_prog_enter+0x9/0x50
       bpf_trampoline_6442466513_0+0x18/0x1000
       migrate_disable+0x5/0x50
       __bpf_prog_enter+0x9/0x50
       bpf_trampoline_6442466513_0+0x18/0x1000
       migrate_disable+0x5/0x50
       __bpf_prog_enter+0x9/0x50
       bpf_trampoline_6442466513_0+0x18/0x1000
       migrate_disable+0x5/0x50
       ...

    Fixing this by adding deny list of btf ids for tracing
    programs and checking btf id during program verification.
    Adding above functions to this list.

    Suggested-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Jiri Olsa <jolsa@kernel.org>
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Link: https://lore.kernel.org/bpf/20210429114712.43783-1-jolsa@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d573fd0394d0e2504989013431cb5bfa2ae6a636
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sat May 8 17:01:46 2021 +0200

    ASoC: Intel: bytcr_rt5640: Add quirk for the Lenovo Miix 3-830 tablet
```

```
    [ Upstream commit f0353e1f53f92f7b3da91e6669f5d58ee222ebe8 ]

    The Lenovo Miix 3-830 tablet has only 1 speaker, has an internal analog
    mic on IN1 and uses JD2 for jack-detect, add a quirk to automatically
    apply these settings on Lenovo Miix 3-830 tablets.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Link: https://lore.kernel.org/r/20210508150146.28403-2-hdegoede@redhat.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 903491d1e16870ffb256f9fa7fcba6d81ee386e2
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sat May 8 17:01:45 2021 +0200

    ASoC: Intel: bytcr_rt5640: Add quirk for the Glavey TM800A550L tablet

    [ Upstream commit 28c268d3acdd4cbcd2ac320b85609e77f84e74a7 ]

    Add a quirk for the Glavey TM800A550L tablet, this BYTCR tablet has no CHAN
    package in its ACPI tables and uses SSP0-AIF1 rather then SSP0-AIF2 which
    is the default for BYTCR devices.

    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Acked-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Link: https://lore.kernel.org/r/20210508150146.28403-1-hdegoede@redhat.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9f1bab36a7090825687ee574317d3c51461e82d7
Author: Dinghao Liu <dinghao.liu@zju.edu.cn>
Date:   Mon Apr 12 13:49:07 2021 +0800

    usb: cdns3: Fix runtime PM imbalance on error

    [ Upstream commit 07adc0225484fc199e3dc15ec889f75f498c4fca ]

    When cdns3_gadget_start() fails, a pairing PM usage counter
    decrement is needed to keep the counter balanced.

    Signed-off-by: Dinghao Liu <dinghao.liu@zju.edu.cn>
    Link: https://lore.kernel.org/r/20210412054908.7975-1-dinghao.liu@zju.edu.cn
    Signed-off-by: Peter Chen <peter.chen@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90d0a3c76965d7a10fc87c07be3e9714e2130d5c
Author: Jeimon <jjjinmeng.zhou@gmail.com>
Date:   Sat May 8 11:52:30 2021 +0800

    net/nfc/rawsock.c: fix a permission check bug

    [ Upstream commit 8ab78863e9eff11910e1ac8bcf478060c29b379e ]

    The function rawsock_create() calls a privileged function sk_alloc(), which requires a
ns-aware check to check net->user_ns, i.e., ns_capable(). However, the original code checks
the init_user_ns using capable(). So we replace the capable() with ns_capable().

    Signed-off-by: Jeimon <jjjinmeng.zhou@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit abcc5436d34167f2c19f8e06ee6f0fe1edf0ef69
Author: Bixuan Cui <cuibixuan@huawei.com>
Date:   Sat May 8 11:15:12 2021 +0800
```

```
    ASoC: codecs: lpass-tx-macro: add missing MODULE_DEVICE_TABLE

    [ Upstream commit 14c0c423746fe7232a093a68809a4bc6233eed60 ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Bixuan Cui <cuibixuan@huawei.com>
    Reviewed-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20210508031512.53783-1-cuibixuan@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cacae854aea33c30ce0bb93c7fd285d91db8d3c
Author: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
Date:   Mon May 10 11:38:44 2021 +0100

    ASoC: codecs: lpass-rx-macro: add missing MODULE_DEVICE_TABLE

    [ Upstream commit d4335d058f8430a0ce2b43dab9531f3a3cf9fe2c ]

    Fix module loading by adding missing MODULE_DEVICE_TABLE.

    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20210510103844.1532-1-srinivas.kandagatla@linaro.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b9cceb297dc5a67597efdb1a0a20138d0025fa7
Author: Jiri Olsa <jolsa@kernel.org>
Date:   Wed May 5 15:25:29 2021 +0200

    bpf: Forbid trampoline attach for functions with variable arguments

    [ Upstream commit 31379397dcc364a59ce764fabb131b645c43e340 ]

    We can't currently allow to attach functions with variable arguments.
    The problem is that we should save all the registers for arguments,
    which is probably doable, but if caller uses more than 6 arguments,
    we need stack data, which will be wrong, because of the extra stack
    frame we do in bpf trampoline, so we could crash.

    Also currently there's malformed trampoline code generated for such
    functions at the moment as described in:

      https://lore.kernel.org/bpf/20210429212834.82621-1-jolsa@kernel.org/

    Signed-off-by: Jiri Olsa <jolsa@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Acked-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20210505132529.401047-1-jolsa@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cbdf358ea88fcc6f9e7e2b1175b6b358b355efd3
Author: Karen Dombroski <karen.dombroski@marsbioimaging.com>
Date:   Wed Apr 28 23:38:02 2021 -0600

    spi: spi-zynq-qspi: Fix stack violation bug

    [ Upstream commit 6d5ff8e632a4f2389c331e5554cd1c2a9a28c7aa ]

    When the number of bytes for the op is greater than one, the read could
    run off the end of the function stack and cause a crash.
```

```
    This patch restores the behaviour of safely reading out of the original
    opcode location.

    Signed-off-by: Karen Dombroski <karen.dombroski@marsbioimaging.com>
    Signed-off-by: Amit Kumar Mahapatra <amit.kumar-mahapatra@xilinx.com>
    Link: https://lore.kernel.org/r/20210429053802.17650-3-amit.kumar-mahapatra@xilinx.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e92ac9263b06c39f05526ad15c90d228cdab60fd
Author: Saravana Kannan <saravanak@google.com>
Date:   Mon Apr 26 16:56:38 2021 -0700

    spi: Fix spi device unregister flow

    [ Upstream commit c7299fea67696db5bd09d924d1f1080d894f92ef ]

    When an SPI device is unregistered, the spi->controller->cleanup() is
    called in the device's release callback. That's wrong for a couple of
    reasons:

    1. spi_dev_put() can be called before spi_add_device() is called. And
       it's spi_add_device() that calls spi_setup(). This will cause clean()
       to get called without the spi device ever being setup.

    2. There's no guarantee that the controller's driver would be present by
       the time the spi device's release function gets called.

    3. It also causes "sleeping in atomic context" stack dump[1] when device
       link deletion code does a put_device() on the spi device.

    Fix these issues by simply moving the cleanup from the device release
    callback to the actual spi_unregister_device() function.

    [1] - https://lore.kernel.org/lkml/CAHp75Vc=FCGcUyS0v6fnxme2YJ+qD+Y-
hQDQLa2JhWNON9VmsQ@mail.gmail.com/

    Signed-off-by: Saravana Kannan <saravanak@google.com>
    Link: https://lore.kernel.org/r/20210426235638.1285530-1-saravanak@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9911e503a80003f8070b689eb8a909800504275f
Author: Vijendar Mukunda <Vijendar.Mukunda@amd.com>
Date:   Wed Apr 28 01:53:31 2021 +0530

    ASoC: amd: fix for pcm_read() error

    [ Upstream commit 6879e8e759bf9e05eaee85e32ca1a936e6b46da1 ]

    Below phython script throwing pcm_read() error.

    import subprocess

    p = subprocess.Popen(["aplay -t raw -D plughw:1,0 /dev/zero"], shell=True)
    subprocess.call(["arecord -Dhw:1,0 --dump-hw-params"], shell=True)
    subprocess.call(["arecord -Dhw:1,0 -fdat -d1 /dev/null"], shell=True)
    p.kill()

    Handling ACP global external interrupt enable register
    causing this issue.
    This register got updated wrongly when there is active
    stream causing interrupts disabled for active stream.
    Refactored code to handle enabling and disabling external interrupts.
```

```
    Signed-off-by: Vijendar Mukunda <Vijendar.Mukunda@amd.com>
    Link: https://lore.kernel.org/r/1619555017-29858-1-git-send-email-
Vijendar.Mukunda@amd.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b6fb1177d5b0b0f108dc8df3db93eaf4e03260b
Author: Marco Felsch <m.felsch@pengutronix.de>
Date:   Fri Apr 23 15:54:02 2021 +0200

    ASoC: max98088: fix ni clock divider calculation

    [ Upstream commit 6c9762a78c325107dc37d20ee21002b841679209 ]

    The ni1/ni2 ratio formula [1] uses the pclk which is the prescaled mclk.
    The max98088 datasheet [2] has no such formula but table-12 equals so
    we can assume that it is the same for both devices.

    While on it make use of DIV_ROUND_CLOSEST_ULL().

    [1] https://datasheets.maximintegrated.com/en/ds/MAX98089.pdf; page 86
    [2] https://datasheets.maximintegrated.com/en/ds/MAX98088.pdf; page 82

    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Link: https://lore.kernel.org/r/20210423135402.32105-1-m.felsch@pengutronix.de
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35b9a92b540819b6ddfd14903e18bf8869909fe3
Author: Kees Cook <keescook@chromium.org>
Date:   Tue Jun 8 10:12:21 2021 -0700

    proc: Track /proc/$pid/attr/ opener mm_struct

    commit 591a22c14d3f45cc38bd1931c593c221df2f1881 upstream.

    Commit bfb819ea20ce ("proc: Check /proc/$pid/attr/ writes against file opener")
    tried to make sure that there could not be a confusion between the opener of
    a /proc/$pid/attr/ file and the writer. It used struct cred to make sure
    the privileges didn't change. However, there were existing cases where a more
    privileged thread was passing the opened fd to a differently privileged thread
    (during container setup). Instead, use mm_struct to track whether the opener
    and writer are still the same process. (This is what several other proc files
    already do, though for different reasons.)

    Reported-by: Christian Brauner <christian.brauner@ubuntu.com>
    Reported-by: Andrea Righi <andrea.righi@canonical.com>
    Tested-by: Andrea Righi <andrea.righi@canonical.com>
    Fixes: bfb819ea20ce ("proc: Check /proc/$pid/attr/ writes against file opener")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```