```
commit 64376a981a0e2e57c46efa63197c2ebb7dab35df
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Tue Jul 20 16:00:42 2021 +0200

    Linux 5.13.4

    Link: https://lore.kernel.org/r/20210719144944.537151528@linuxfoundation.org
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Link: https://lore.kernel.org/r/20210719184345.989046417@linuxfoundation.org
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 71de462034c69525a5049fbdf3903c5833cbce04
Author: Eric Sandeen <sandeen@redhat.com>
Date:   Tue Jul 13 17:49:23 2021 +0200

    seq_file: disallow extremely large seq buffer allocations

    commit 8cae8cd89f05f6de223d63e6d15e31c8ba9cf53b upstream.

    There is no reasonable need for a buffer larger than this, and it avoids
    int overflow pitfalls.

    Fixes: 058504edd026 ("fs/seq_file: fallback to vmalloc allocation")
    Suggested-by: Al Viro <viro@zeniv.linux.org.uk>
    Reported-by: Qualys Security Advisory <qsa@qualys.com>
    Signed-off-by: Eric Sandeen <sandeen@redhat.com>
    Cc: stable@kernel.org
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 31cf8e7e504eb91f697bb19c4175a53e5cbb4b2c
Author: Tong Zhang <ztong0001@gmail.com>
Date:   Sat May 22 00:37:25 2021 -0400

    misc: alcor_pci: fix inverted branch condition

    commit 281e468446994a7672733af2bf941f4110d4a895 upstream.

    This patch fixes a trivial mistake that I made in the previous attempt
    in fixing the null bridge issue. The branch condition is inverted and we
    should call alcor_pci_find_cap_offset() only if bridge is not null.

    Reported-by: Colin Ian King <colin.king@canonical.com>
    Fixes: 3ce3e45cc333 ("misc: alcor_pci: fix null-ptr-deref when there is no PCI bridge")
    Signed-off-by: Tong Zhang <ztong0001@gmail.com>
    Link: https://lore.kernel.org/r/20210522043725.602179-1-ztong0001@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9930eeb4ffae7c0922f7b0539dfe141763f0262c
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Thu Jun 3 15:33:20 2021 +0300

    scsi: scsi_dh_alua: Fix signedness bug in alua_rtpg()

    commit 80927822e8b6be46f488524cd7d5fe683de97fc4 upstream.

    The "retval" variable needs to be signed for the error handling to work.

    Link: https://lore.kernel.org/r/YLjMEAFNxOas1mIp@mwanda
    Fixes: 7e26e3ea0287 ("scsi: scsi_dh_alua: Check for negative result value")
    Reviewed-by: Martin Wilck <mwilck@suse.com>
```

```
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e1b2b2b61d30d7ce057ec17237c217d152ed97f2
Author: Viresh Kumar <viresh.kumar@linaro.org>
Date:   Fri Jun 18 13:31:27 2021 +0530

    cpufreq: CPPC: Fix potential memleak in cppc_cpufreq_cpu_init

    [ Upstream commit fe2535a44904a77615a3af8e8fd7dafb98fb0e1b ]

    It's a classic example of memleak, we allocate something, we fail and
    never free the resources.

    Make sure we free all resources on policy ->init() failures.

    Fixes: a28b2bfc099c ("cppc_cpufreq: replace per-cpu data array with a list")
    Tested-by: Vincent Guittot <vincent.guittot@linaro.org>
    Reviewed-by: Ionela Voinescu <ionela.voinescu@arm.com>
    Tested-by: Qian Cai <quic_qiancai@quicinc.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c898c9bfd3d58d5ef7ab1dd5647df5fdb39ac9d7
Author: Jin Yao <yao.jin@linux.intel.com>
Date:   Thu Jul 1 14:42:53 2021 +0800

    perf tools: Fix pattern matching for same substring in different PMU type

    [ Upstream commit c47a5599eda324bacdacd125227a0925d6c50fbe ]

    Some different PMU types may have the same substring. For example, on
    Icelake server we have PMU types "uncore_imc" and
    "uncore_imc_free_running". Both PMU types have the substring
    "uncore_imc".  But the parser wrongly thinks they are the same PMU type.

    We enable an imc event,
    perf stat -e uncore_imc/event=0xe3/ -a -- sleep 1

    Perf actually expands the event to:

      uncore_imc_0/event=0xe3/
      uncore_imc_1/event=0xe3/
      uncore_imc_2/event=0xe3/
      uncore_imc_3/event=0xe3/
      uncore_imc_4/event=0xe3/
      uncore_imc_5/event=0xe3/
      uncore_imc_6/event=0xe3/
      uncore_imc_7/event=0xe3/
      uncore_imc_free_running_0/event=0xe3/
      uncore_imc_free_running_1/event=0xe3/
      uncore_imc_free_running_3/event=0xe3/
      uncore_imc_free_running_4/event=0xe3/

    That's because the "uncore_imc_free_running" matches the
    pattern "uncore_imc*".

    Now we check that the last characters of PMU name is '_<digit>'.

    For example, for pattern "uncore_imc*", "uncore_imc_0" is parsed ok, but
    "uncore_imc_free_running_0" fails.

    Fixes: b2b9d3a3f0211c5d ("perf pmu: Support wildcards on pmu name in dynamic pmu
events")
```

```
    Signed-off-by: Jin Yao <yao.jin@linux.intel.com>
    Reviewed-by: Kan Liang <kan.liang@linux.intel.com>
    Acked-by: Jiri Olsa <jolsa@redhat.com>
    Cc: Agustin Vega-Frias <agustinv@codeaurora.org>
    Cc: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Cc: Andi Kleen <ak@linux.intel.com>
    Cc: Peter Zijlstra <peterz@infradead.org>
    Link: http://lore.kernel.org/lkml/20210701064253.1175-1-yao.jin@linux.intel.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04ea267f5357ddcf39ff7098cdd03c855b9fa036
Author: Martin Fäcknitz <faecknitz@hotsplots.de>
Date:   Mon Jul 5 02:03:54 2021 +0200

    MIPS: vdso: Invalid GIC access through VDSO

    [ Upstream commit 47ce8527fbba145a7723685bc9a27d9855e06491 ]

    Accessing raw timers (currently only CLOCK_MONOTONIC_RAW) through VDSO
    doesn't return the correct time when using the GIC as clock source.
    The address of the GIC mapped page is in this case not calculated
    correctly. The GIC mapped page is calculated from the VDSO data by
    subtracting PAGE_SIZE:

      void *get_gic(const struct vdso_data *data) {
        return (void __iomem *)data - PAGE_SIZE;
      }

    However, the data pointer is not page aligned for raw clock sources.
    This is because the VDSO data for raw clock sources (CS_RAW = 1) is
    stored after the VDSO data for coarse clock sources (CS_HRES_COARSE = 0).
    Therefore, only the VDSO data for CS_HRES_COARSE is page aligned:

      +--------------------+
      |                    |
      | vd[CS_RAW]         | ---+
      | vd[CS_HRES_COARSE] |    |
      +--------------------+    | -PAGE_SIZE
      |                    |    |
      |  GIC mapped page   | <--+
      |                    |
      +--------------------+

    When __arch_get_hw_counter() is called with &vd[CS_RAW], get_gic returns
    the wrong address (somewhere inside the GIC mapped page). The GIC counter
    values are not returned which results in an invalid time.

    Fixes: a7f4df4e21dd ("MIPS: VDSO: Add implementations of gettimeofday() and
clock_gettime()")
    Signed-off-by: Martin Fäcknitz <faecknitz@hotsplots.de>
    Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bd5b2e8f4a7ce81b874096cef84bf17270d2af94
Author: Sven Schnelle <svens@linux.ibm.com>
Date:   Tue Jul 6 20:24:58 2021 +0200

    s390/irq: remove HAVE_IRQ_EXIT_ON_IRQ_STACK

    [ Upstream commit 0aa4ff7688632a86bdb133fa106f2ccd514b91a7 ]

    This is no longer true since we switched to generic entry. The code
    switches to the IRQ stack before calling do_IRQ, but switches back
    to the kernel stack before calling irq_exit().
```

```
    Fixes: 56e62a737028 ("s390: convert to generic entry")
    Signed-off-by: Sven Schnelle <svens@linux.ibm.com>
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 272fa75ec8a00be286ca34fbdec90682a4a55dab
Author: Kajol Jain <kjain@linux.ibm.com>
Date:   Mon Jun 28 11:53:41 2021 +0530

    perf script python: Fix buffer size to report iregs in perf script

    [ Upstream commit dea8cfcc33695f70f56023b416cf88ae44c8a45a ]

    Commit 48a1f565261d2ab1 ("perf script python: Add more PMU fields to
    event handler dict") added functionality to report fields like weight,
    iregs, uregs etc via perf report.  That commit predefined buffer size to
    512 bytes to print those fields.

    But in PowerPC, since we added extended regs support in:

      068aeea3773a6f4c ("perf powerpc: Support exposing Performance Monitor Counter SPRs as
part of extended regs")
      d735599a069f6936 ("powerpc/perf: Add extended regs support for power10 platform")

    Now iregs can carry more bytes of data and this predefined buffer size
    can result to data loss in perf script output.

    This patch resolves this issue by making the buffer size dynamic, based
    on the number of registers needed to print. It also changes the
    regs_map() return type from int to void, as it is not being used by the
    set_regs_in_dict(), its only caller.

    Fixes: 068aeea3773a6f4c ("perf powerpc: Support exposing Performance Monitor Counter
SPRs as part of extended regs")
    Signed-off-by: Kajol Jain <kjain@linux.ibm.com>
    Tested-by: Nageswara R Sastry <rnsastry@linux.ibm.com>
    Cc: Athira Jajeev <atrajeev@linux.vnet.ibm.com>
    Cc: Jiri Olsa <jolsa@redhat.com>
    Cc: Madhavan Srinivasan <maddy@linux.vnet.ibm.com>
    Cc: Paul Clarke <pc@us.ibm.com>
    Cc: Ravi Bangoria <ravi.bangoria@linux.ibm.com>
    Cc: linuxppc-dev@lists.ozlabs.org
    Link: http://lore.kernel.org/lkml/20210628062341.155839-1-kjain@linux.ibm.com
    Signed-off-by: Arnaldo Carvalho de Melo <acme@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3a5f4fb0ee2d8b4f4bf456a715ba33977851529e
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Jul 4 16:02:11 2021 -0700

    mips: disable branch profiling in boot/decompress.o

    [ Upstream commit 97e488073cfca0eea84450169ca4cbfcc64e33e3 ]

    Use DISABLE_BRANCH_PROFILING for arch/mips/boot/compressed/decompress.o
    to prevent linkage errors.

    mips64-linux-ld: arch/mips/boot/compressed/decompress.o: in function
`LZ4_decompress_fast_extDict':
    decompress.c:(.text+0x8c): undefined reference to `ftrace_likely_update'
    mips64-linux-ld: decompress.c:(.text+0xf4): undefined reference to
`ftrace_likely_update'
    mips64-linux-ld: decompress.c:(.text+0x200): undefined reference to
```

```
  `ftrace_likely_update'
      mips64-linux-ld: decompress.c:(.text+0x230): undefined reference to
  `ftrace_likely_update'
      mips64-linux-ld: decompress.c:(.text+0x320): undefined reference to
  `ftrace_likely_update'
      mips64-linux-ld: arch/mips/boot/compressed/decompress.o:decompress.c:(.text+0x3f4): more
  undefined references to `ftrace_likely_update' follow

      Fixes: e76e1fdfa8f8 ("lib: add support for LZ4-compressed kernel")
      Reported-by: kernel test robot <lkp@intel.com>
      Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
      Cc: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
      Cc: linux-mips@vger.kernel.org
      Cc: Kyungsik Lee <kyungsik.lee@lge.com>
      Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 79a0e94e1434f3e1df5f053599c4953d610971d2
  Author: Arnd Bergmann <arnd@arndb.de>
  Date:   Fri Jul 2 16:28:37 2021 +0200

      mips: always link byteswap helpers into decompressor

      [ Upstream commit cddc40f5617e53f97ef019d5b29c1bd6cbb031ec ]

      My series to clean up the unaligned access implementation
      across architectures caused some mips randconfig builds to
      fail with:

         mips64-linux-ld: arch/mips/boot/compressed/decompress.o: in function
  `decompress_kernel':
         decompress.c:(.text.decompress_kernel+0x54): undefined reference to `__bswapsi2'

      It turns out that this problem has already been fixed for the XZ
      decompressor but now it also shows up in (at least) LZO and LZ4.  From my
      analysis I concluded that the compiler could always have emitted those
      calls, but the different implementation allowed it to make otherwise
      better decisions about not inlining the byteswap, which results in the
      link error when the out-of-line code is missing.

      While it could be addressed by adding it to the two decompressor
      implementations that are known to be affected, but as this only adds
      112 bytes to the kernel, the safer choice is to always add them.

      Fixes: c50ec6787536 ("MIPS: zboot: Fix the build with XZ compression on older GCC
  versions")
      Fixes: 0652035a5794 ("asm-generic: unaligned: remove byteshift helpers")
      Link: https://lore.kernel.org/linux-mm/202106301304.gz2wVY9w-lkp@intel.com/
      Link: https://lore.kernel.org/linux-mm/202106260659.TyMe8mjr-lkp@intel.com/
      Link: https://lore.kernel.org/linux-mm/202106172016.onWT6Tza-lkp@intel.com/
      Link: https://lore.kernel.org/linux-mm/202105231743.JJcALnhS-lkp@intel.com/
      Signed-off-by: Arnd Bergmann <arnd@arndb.de>
      Signed-off-by: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 8be9eadd7bd646a0dd35754f5f64a57ee3bb544b
  Author: Peter Zijlstra <peterz@infradead.org>
  Date:   Mon Jun 28 13:24:12 2021 +0200

      kprobe/static_call: Restore missing static_call_text_reserved()

      [ Upstream commit fa68bd09fc62240a383c0c601d3349c47db10c34 ]

      Restore two hunks from commit:
```

```
       6333e8f73b83 ("static_call: Avoid kprobes on inline static_call()s")

    that went walkabout in a Git merge commit.

    Fixes: 76d4acf22b48 ("Merge tag 'perf-kprobes-2020-12-14' of
git://git.kernel.org/pub/scm/linux/kernel/git/tip/tip")
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Acked-by: Masami Hiramatsu <mhiramat@kernel.org>
    Link: https://lore.kernel.org/r/20210628113045.167127609@infradead.org
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fda47dd9f5d7a026f4eae3cc249b69f695df2182
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Mon Jun 28 13:24:11 2021 +0200

    static_call: Fix static_call_text_reserved() vs __init

    [ Upstream commit 2bee6d16e4379326b1eea454e68c98b17456769e ]

    It turns out that static_call_text_reserved() was reporting __init
    text as being reserved past the time when the __init text was freed
    and re-used.

    This is mostly harmless and will at worst result in refusing a kprobe.

    Fixes: 6333e8f73b83 ("static_call: Avoid kprobes on inline static_call()s")
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Reviewed-by: Masami Hiramatsu <mhiramat@kernel.org>
    Link: https://lore.kernel.org/r/20210628113045.106211657@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a861b41d0b41fd52a8fbc408f54a04ca795d2f30
Author: Peter Zijlstra <peterz@infradead.org>
Date:   Mon Jun 28 13:24:10 2021 +0200

    jump_label: Fix jump_label_text_reserved() vs __init

    [ Upstream commit 9e667624c291753b8a5128f620f493d0b5226063 ]

    It turns out that jump_label_text_reserved() was reporting __init text
    as being reserved past the time when the __init text was freed and
    re-used.

    For a long time, this resulted in, at worst, not being able to kprobe
    text that happened to land at the re-used address. However a recent
    commit e7bf1ba97afd ("jump_label, x86: Emit short JMP") made it a
    fatal mistake because it now needs to read the instruction in order to
    determine the conflict -- an instruction that's no longer there.

    Fixes: 4c3ef6d79328 ("jump label: Add jump_label_text_reserved() to reserve jump
points")
    Reported-by: kernel test robot <oliver.sang@intel.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Reviewed-by: Masami Hiramatsu <mhiramat@kernel.org>
    Link: https://lore.kernel.org/r/20210628113045.045141693@infradead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7ec843ba963e6f1f85c56ed2f00cbb2c0f30aee6
Author: Xuewen Yan <xuewen.yan@unisoc.com>
Date:   Wed Jun 30 22:12:04 2021 +0800

    sched/uclamp: Ignore max aggregation if rq is idle
```

```
    [ Upstream commit 3e1493f46390618ea78607cb30c58fc19e2a5035 ]

    When a task wakes up on an idle rq, uclamp_rq_util_with() would max
    aggregate with rq value. But since there is no task enqueued yet, the
    values are stale based on the last task that was running. When the new
    task actually wakes up and enqueued, then the rq uclamp values should
    reflect that of the newly woken up task effective uclamp values.

    This is a problem particularly for uclamp_max because it default to
    1024. If a task p with uclamp_max = 512 wakes up, then max aggregation
    would ignore the capping that should apply when this task is enqueued,
    which is wrong.

    Fix that by ignoring max aggregation if the rq is idle since in that
    case the effective uclamp value of the rq will be the ones of the task
    that will wake up.

    Fixes: 9d20ad7dfc9a ("sched/uclamp: Add uclamp_util_with()")
    Signed-off-by: Xuewen Yan <xuewen.yan@unisoc.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Valentin Schneider <valentin.schneider@arm.com>
    [qias: Changelog]
    Reviewed-by: Qais Yousef <qais.yousef@arm.com>
    Link: https://lore.kernel.org/r/20210630141204.8197-1-xuewen.yan94@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c368aa2bbd2ab028b69ec9933d2494f12881ae2d
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Sat Jun 12 09:18:34 2021 +0200

    scsi: be2iscsi: Fix an error handling path in beiscsi_dev_probe()

    [ Upstream commit 030e4138d11fced3b831c2761e4cecf347bae99c ]

    If an error occurs after a pci_enable_pcie_error_reporting() call, it must
    be undone by a corresponding pci_disable_pcie_error_reporting() call, as
    already done in the remove function.

    Link:
https://lore.kernel.org/r/77adb02cfea7f1364e5603ecf3930d8597ae356e.1623482155.git.christophe
.jaillet@wanadoo.fr
    Fixes: 3567f36a09d1 ("[SCSI] be2iscsi: Fix AER handling in driver")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6a08c2f1c413abde5e725337bd28b18d0cbc4ca
Author: Alex Bee <knaerzche@gmail.com>
Date:   Sat Jun 19 14:13:06 2021 +0200

    arm64: dts: rockchip: Re-add regulator-always-on for vcc_sdio for rk3399-roc-pc

    [ Upstream commit eb607cd4957fb0ef97beb2a8293478be6a54240a ]

    Re-add the regulator-always-on property for vcc_sdio which supplies sdmmc,
    since it gets disabled during reboot now and the bootrom expects it to be
    enabled  when booting from SD card. This makes rebooting impossible in that
    case and requires a hard reset to boot again.

    Fixes: 04a0077fdb19 ("arm64: dts: rockchip: Remove always-on properties from regulator
nodes on rk3399-roc-pc.")
    Signed-off-by: Alex Bee <knaerzche@gmail.com>
    Link: https://lore.kernel.org/r/20210619121306.7740-1-knaerzche@gmail.com
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
```

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

commit 12f8ee513a33b5de27a29538ca267ce0129ab55f
Author: Alex Bee <knaerzche@gmail.com>
Date:   Sat Jun 19 14:14:46 2021 +0200

    arm64: dts: rockchip: Re-add regulator-boot-on, regulator-always-on for vdd_gpu on
rk3399-roc-pc

    [ Upstream commit 06b2818678d9b35102c9816ffaf6893caf306ed0 ]

    This might be a limitation of either the current panfrost driver
    devfreq implementation or how the gpu is implemented in RK3399 SoC.
    The gpu regulator must never get disabled or the registers get
    (randomly?) inaccessable by the driver. (see all other RK3399 boards)

    Fixes: ec7d731d81e7 ("arm64: dts: rockchip: Add node for gpu on rk3399-roc-pc")
    Signed-off-by: Alex Bee <knaerzche@gmail.com>
    Link: https://lore.kernel.org/r/20210619121446.7802-1-knaerzche@gmail.com
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9f5d2c5052d3a413f995bdcb944c3b1ab6d5b9cf
Author: Pali Rohár <pali@kernel.org>
Date:   Thu May 20 13:35:20 2021 +0200

    firmware: turris-mox-rwtm: show message about HWRNG registration

    [ Upstream commit fae2016099226943150 7708fb74c1fd9f3c309c1 ]

    Currently it is hard to determinate if on Armada 3720 device is HWRNG
    by running kernel accessible or not. So print information message into
    dmesg when HWRNG is available and registration was successful.

    Fixes: 389711b37493 ("firmware: Add Turris Mox rWTM firmware driver")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 852ae2d33e1fecac115d4ed12091a5703c989a3e
Author: Pali Rohár <pali@kernel.org>
Date:   Thu May 20 13:35:19 2021 +0200

    firmware: turris-mox-rwtm: fail probing when firmware does not support hwrng

    [ Upstream commit 2eab59cf0d2036a5a9e264f719b71c21ccf679c2 ]

    When Marvell's rWTM firmware, which does not support the GET_RANDOM
    command, is used, kernel prints an error message
      hwrng: no data available
    every 10 seconds.

    Fail probing of this driver if the rWTM firmware does not support the
    GET_RANDOM command.

    Fixes: 389711b37493 ("firmware: Add Turris Mox rWTM firmware driver")
    Signed-off-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f6ffd65b7dcfe69a9e64463bfff832e8b932b8a
Author: Marek Behún <kabel@kernel.org>
```

```
Date:   Thu May 20 13:35:18 2021 +0200

    firmware: turris-mox-rwtm: report failures better

    [ Upstream commit 72f99888944c44de1c899bbe44db1e53bdc9d994 ]

    Report a notice level message if a command is not supported by the rWTM
    firmware.

    This should not be an error, merely a notice, because the firmware can
    be used on boards that do not have manufacturing information burned.

    Fixes: 389711b37493 ("firmware: Add Turris Mox rWTM firmware driver")
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Reviewed-by: Pali Rohár <pali@kernel.org>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa1b734129a750d3352befda2141f49d562112ff
Author: Marek Behún <kabel@kernel.org>
Date:   Thu May 20 13:35:17 2021 +0200

    firmware: turris-mox-rwtm: fix reply status decoding function

    [ Upstream commit e34e60253d9272311831daed8a2d967cf80ca3dc ]

    The status decoding function mox_get_status() currently contains an
    incorrect check: if the error status is not MBOX_STS_SUCCESS, it always
    returns -EIO, so the comparison to MBOX_STS_FAIL is never executed and
    we don't get the actual error code sent by the firmware.

    Fix this.

    Signed-off-by: Marek Behún <kabel@kernel.org>
    Reviewed-by: Pali Rohár <pali@kernel.org>
    Reviewed-by: Andrew Lunn <andrew@lunn.ch>
    Fixes: 389711b37493 ("firmware: Add Turris Mox rWTM firmware driver")
    Signed-off-by: Gregory CLEMENT <gregory.clement@bootlin.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3dbb2de9e7afe0005b8c8f9e0da9f13a80e386d7
Author: Masahiro Yamada <masahiroy@kernel.org>
Date:   Wed Jun 2 23:02:13 2021 +0900

    kbuild: remove trailing slashes from $(KBUILD_EXTMOD)

    [ Upstream commit 74ee585b7eecd98be3650e677625a0ee588d08e0 ]

    M= (or KBUILD_EXTMOD) generally expects a directory path without any
    trailing slashes, like M=a/b/c.

    If you add a trailing slash, like M=a/b/c/, you will get ugly build
    logs (two slashes in a series), but it still works fine as long as it
    is consistent between 'make modules' and 'make modules_install'.

    The following commands correctly build and install the modules.

      $ make M=a/b/c/ modules
      $ sudo make M=a/b/c/ modules_install

    Since commit ccae4cfa7bfb ("kbuild: refactor scripts/Makefile.modinst"),
    a problem happens if you add a trailing slash only for modules_install.

      $ make M=a/b/c modules
```

```
       $ sudo make M=a/b/c/ modules_install

    No module is installed in this case, Johannes Berg reported. [1]

    Trim any trailing slashes from $(KBUILD_EXTMOD).

    I used the 'dirname' command to remove all the trailing slashes in
    case someone adds more slashes like M=a/b/c/////. The Make's built-in
    function, $(dir ...) cannot take care of such a case.

    [1]:
https://lore.kernel.org/lkml/10cc8522b27a051e6a9c3e158a4c4b6414fd04a0.camel@sipsolutions.net
/

    Fixes: ccae4cfa7bfb ("kbuild: refactor scripts/Makefile.modinst")
    Reported-by: Johannes Berg <johannes@sipsolutions.net>
    Signed-off-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2292bba57bb6c67ecc5223a616d809519eabac80
Author: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
Date:   Sat Jun 5 10:52:11 2021 +0200

    thermal/drivers/rcar_gen3_thermal: Fix coefficient calculations

    [ Upstream commit 8946187ab57ffd02088e50256c73dd31f49db06d ]

    The fixed value of 157 used in the calculations are only correct for
    M3-W, on other Gen3 SoC it should be 167. The constant can be derived
    correctly from the static TJ_3 constant and the SoC specific TJ_1 value.
    Update the calculation be correct on all Gen3 SoCs.

    Fixes: 4eb39f79ef44 ("thermal: rcar_gen3_thermal: Update value of Tj_1")
    Reported-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Signed-off-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Reviewed-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20210605085211.564909-1-
niklas.soderlund+renesas@ragnatech.se
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 31402780830329bb8049eafcb023b6ebcf16e5ab
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Tue Jun 8 10:44:14 2021 +0530

    arm64: dts: ti: k3-am642-evm: align ti,pindir-d0-out-d1-in property with dt-shema

    [ Upstream commit d3f1b155c04d949c843e6028034766aba1e0f8bf ]

    ti,pindir-d0-out-d1-in property is expected to be of type boolean.
    Therefore, fix the property accordingly.

    Fixes: 4fb6c04683aa ("arm64: dts: ti: k3-am642-evm: Add support for SPI EEPROM")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Reviewed-by: Vignesh Raghavendra <vigneshr@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210608051414.14873-3-a-govindraju@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 908b140229a206dafc659b620144351c982a9393
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Tue Jun 8 10:44:13 2021 +0530

    arm64: dts: ti: am65: align ti,pindir-d0-out-d1-in property with dt-shema
```

```
    [ Upstream commit 4f76ea7b4da1cce9a9bda1fa678ef8036584c66b ]

    ti,pindir-d0-out-d1-in property is expected to be of type boolean.
    Therefore, fix the property accordingly.

    Fixes: e180f76d0641 ("arm64: dts: ti: Add support for Siemens IOT2050 boards")
    Fixes: 5da94b50475a ("arm64: dts: ti: k3-am654: Enable main domain McSPI0")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Acked-by: Jan Kiszka <jan.kiszka@siemens.com>
    Reviewed-by: Vignesh Raghavendra <vigneshr@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210608051414.14873-2-a-govindraju@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6f62efde1ae09b22ebbc1a80647aea76ff06328d
Author: Grygorii Strashko <grygorii.strashko@ti.com>
Date:   Tue Jun 8 21:49:40 2021 +0300

    arm64: dts: ti: k3-am642-main: fix ports mac properties

    [ Upstream commit 50c9bfca1bfe9ffd56d8c5deecf9204d14e20bfd ]

    The current device tree CPSW3g node adds non-zero "mac-address" property to
    the ports, which prevents random MAC address assignment to network devices
    if bootloader failed to update DT. This may cause more then one host to
    have the same MAC in the network.

     mac-address = [00 00 de ad be ef];
     mac-address = [00 01 de ad be ef];

    In addition, there is one MAC address available in eFuse registers which
    can be used for default port 1.

    Hence, fix ports MAC properties by:
    - resetting "mac-address" property to 0
    - adding ti,syscon-efuse = <&main_conf 0x200> to Port 1

    Fixes: 3753b12877b6 ("arm64: dts: ti: k3-am64-main: Add CPSW DT node")
    Signed-off-by: Grygorii Strashko <grygorii.strashko@ti.com>
    Reviewed-by: Vignesh Raghavendra <vigneshr@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210608184940.25934-1-grygorii.strashko@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b47aec030e9da5b4fc140d3f1504d782df7aea0
Author: Christoph Niedermaier <cniedermaier@dh-electronics.com>
Date:   Wed May 26 12:54:00 2021 +0200

    ARM: dts: imx6q-dhcom: Add gpios pinctrl for i2c bus recovery

    [ Upstream commit ddc873cd3c0af4faad6a00bffda21c3f775126dd ]

    The i2c bus can freeze at the end of transaction so the bus can no longer work.
    This scenario is improved by adding scl/sda gpios definitions to implement the
    i2c bus recovery mechanism.

    Fixes: 52c7a088badd ("ARM: dts: imx6q: Add support for the DHCOM iMX6 SoM and PDK2")
    Signed-off-by: Christoph Niedermaier <cniedermaier@dh-electronics.com>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Cc: Fabio Estevam <festevam@gmail.com>
    Cc: Marek Vasut <marex@denx.de>
    Cc: NXP Linux Team <linux-imx@nxp.com>
    Cc: kernel@dh-electronics.com
    To: linux-arm-kernel@lists.infradead.org
```

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

commit 189c0feed615e66b72100a4b58cfead97d204afb
Author: Christoph Niedermaier <cniedermaier@dh-electronics.com>
Date:   Wed May 26 12:53:59 2021 +0200

    ARM: dts: imx6q-dhcom: Fix ethernet plugin detection problems

    [ Upstream commit e2bdd3484890441b9cc2560413a86e8f2aa04157 ]

    To make the ethernet cable plugin detection reliable the
    power detection of the smsc phy has been disabled.

    Fixes: 52c7a088badd ("ARM: dts: imx6q: Add support for the DHCOM iMX6 SoM and PDK2")
    Signed-off-by: Christoph Niedermaier <cniedermaier@dh-electronics.com>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Cc: Fabio Estevam <festevam@gmail.com>
    Cc: Marek Vasut <marex@denx.de>
    Cc: NXP Linux Team <linux-imx@nxp.com>
    Cc: kernel@dh-electronics.com
    To: linux-arm-kernel@lists.infradead.org
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bffed302186c0707cdc4c021c194fce8c82d97bf
Author: Christoph Niedermaier <cniedermaier@dh-electronics.com>
Date:   Wed May 26 12:53:58 2021 +0200

    ARM: dts: imx6q-dhcom: Fix ethernet reset time properties

    [ Upstream commit c016c26c1631f539c652b5d82242a3ca402545c1 ]

    Fix ethernet reset time properties as described in
    Documentation/devicetree/bindings/net/ethernet-phy.yaml

    Fixes: 52c7a088badd ("ARM: dts: imx6q: Add support for the DHCOM iMX6 SoM and PDK2")
    Signed-off-by: Christoph Niedermaier <cniedermaier@dh-electronics.com>
    Cc: Shawn Guo <shawnguo@kernel.org>
    Cc: Fabio Estevam <festevam@gmail.com>
    Cc: Marek Vasut <marex@denx.de>
    Cc: NXP Linux Team <linux-imx@nxp.com>
    Cc: kernel@dh-electronics.com
    To: linux-arm-kernel@lists.infradead.org
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 527b863b5260e26dbd6796bd01ee786a90cf63d7
Author: Chunyan Zhang <chunyan.zhang@unisoc.com>
Date:   Wed May 12 17:37:52 2021 +0800

    thermal/drivers/sprd: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 4d57fd9aeaa013a245bf1fade81e2c30a5efd491 ]

    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.

    Fixes: 554fdbaf19b18 ("thermal: sprd: Add Spreadtrum thermal driver support")
    Signed-off-by: Chunyan Zhang <chunyan.zhang@unisoc.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20210512093752.243168-1-zhang.lyra@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 68812676606765b3e50aa83093640f850f288703
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Tue Jun 8 10:39:52 2021 +0530

    ARM: dts: am437x: align ti,pindir-d0-out-d1-in property with dt-shema

    [ Upstream commit 9b11fec7345f21995f4ea4bafb0e108b9a620238 ]

    ti,pindir-d0-out-d1-in property is expected to be of type boolean.
    Therefore, fix the property accordingly.

    Fixes: b0b039515445 ("ARM: dts: am43x-epos-evm: set data pin directions for spi0 and
spi1")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ca219ca56e48b560e735019662468ce08c2981b3
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Tue Jun 8 10:39:51 2021 +0530

    ARM: dts: am335x: align ti,pindir-d0-out-d1-in property with dt-shema

    [ Upstream commit 414bfe1d26b60ef20b58e36efd5363188a694bab ]

    ti,pindir-d0-out-d1-in property is expected to be of type boolean.
    Therefore, fix the property accordingly.

    Fixes: 444d66fafab8 ("ARM: dts: add spi wifi support to cm-t335")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 75e9d92befadbb5c8f2ca9efd90bcc5aa8483ee9
Author: Gowtham Tammana <g-tammana@ti.com>
Date:   Wed Jun 2 17:04:58 2021 -0500

    ARM: dts: dra7: Fix duplicate USB4 target module node

    [ Upstream commit 78b4b165280d3d70e7a217599f0c06a4c0bb11f9 ]

    With [1] USB4 target-module node got defined in dra74x.dtsi file.
    However, the earlier definition in [2] was not removed, and this
    duplication of the target module is causing boot failure on dra74
    variant boards - dra7-evm, dra76-evm.

    USB4 is only present in DRA74x variants, so keeping the entry in
    dra74x.dtsi and removing it from the top level interconnect hierarchy
    dra7-l4.dtsi file. This change makes the USB4 target module no longer
    visible to AM5718, DRA71x and DRA72x so removing references to it in
    their respective dts files.

    [1]: commit c7b72abca61ec ("ARM: OMAP2+: Drop legacy platform data for
    dra7 dwc3")
    [2]: commit 549fce068a311 ("ARM: dts: dra7: Add l4 interconnect
    hierarchy and ti-sysc data")

    Fixes: c7b72abca61ec ("ARM: OMAP2+: Drop legacy platform data for dra7 dwc3")
    Signed-off-by: Gowtham Tammana <g-tammana@ti.com>
    Reviewed-by: Grygorii Strashko <grygorii.strashko@ti.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4264bda03ef1d9026dffbb9b425646b3f6fba72e
```

```
Author: Icenowy Zheng <icenowy@aosc.io>
Date:   Wed Jun 9 16:38:43 2021 +0800

    arm64: dts: allwinner: a64-sopine-baseboard: change RGMII mode to TXID

    [ Upstream commit bd5431b2f9b30a70f6ed964dd5ee9a6d1c397c06 ]

    Although the schematics of Pine A64-LTS and SoPine Baseboard shows both
    the RX and TX internal delay are enabled, they're using the same broken
    RTL8211E chip batch with Pine A64+, so they should use TXID instead, not
    ID.

    In addition, by checking the real components soldered on both a SoPine
    Baseboard and a Pine A64-LTS, RX delay is not enabled (GR69 soldered and
    GR70 NC) despite the schematics says it's enabled. It's a common
    situation for Pine64 boards that the NC information on schematics is not
    the same with the board.

    So the RGMII delay mode should be TXID on these boards.

    Fixes: c2b111e59a7b ("arm64: dts: allwinner: A64 Sopine: phy-mode rgmii-id")
    Signed-off-by: Icenowy Zheng <icenowy@aosc.io>
    Signed-off-by: Maxime Ripard <maxime@cerno.tech>
    Link: https://lore.kernel.org/r/20210609083843.463750-1-icenowy@aosc.io
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6b45b4932f7b0c36b41fb56a35ad679ece939a0
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Thu May 27 11:43:22 2021 -0400

    memory: fsl_ifc: fix leak of private memory on probe failure

    [ Upstream commit 8e0d09b1232d0538066c40ed4c13086faccbdff6 ]

    On probe error the driver should free the memory allocated for private
    structure.  Fix this by using resource-managed allocation.

    Fixes: a20cbdeffce2 ("powerpc/fsl: Add support for Integrated Flash Controller")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210527154322.81253-2-krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 28f71fd81ebd3b386bf5c7c5539664156f7d72c1
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Thu May 27 11:43:21 2021 -0400

    memory: fsl_ifc: fix leak of IO mapping on probe failure

    [ Upstream commit 3b132ab67fc7a358fff35e808fa65d4bea452521 ]

    On probe error the driver should unmap the IO memory.  Smatch reports:

      drivers/memory/fsl_ifc.c:298 fsl_ifc_ctrl_probe() warn: 'fsl_ifc_ctrl_dev->gregs' not
released on lines: 298.

    Fixes: a20cbdeffce2 ("powerpc/fsl: Add support for Integrated Flash Controller")
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210527154322.81253-1-krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5efc32c0cf6ddefa7bb0cc4e4a49e1f14ba44dbc
Author: Kishon Vijay Abraham I <kishon@ti.com>
Date:   Thu Jun 3 20:04:27 2021 +0530
```

```
    arm64: dts: ti: k3-j721e-common-proc-board: Re-name "link" name as "phy"

    [ Upstream commit 02b4d9186121d842a53e347f53a86ec7f2c6b0c7 ]

    Commit 66db854b1f62d ("arm64: dts: ti: k3-j721e-common-proc-board:
    Configure the PCIe instances") and
    commit 02c35dca2b488 ("arm64: dts: ti: k3-j721e: Enable Super-Speed
    support for USB0") added PHY DT nodes with node name as "link"
    However nodes with #phy-cells should be named 'phy' as discussed in [1].
    Re-name subnodes of serdes in J721E to 'phy'.

    [1] -> http://lore.kernel.org/r/20200909203631.GA3026331@bogus

    Fixes: 66db854b1f62d ("arm64: dts: ti: k3-j721e-common-proc-board: Configure the PCIe
instances")
    Fixes: 02c35dca2b488 ("arm64: dts: ti: k3-j721e: Enable Super-Speed support for USB0")
    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210603143427.28735-5-kishon@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c854fab6a0c94341fe93dbfa55bfe3594a322bfd
Author: Kishon Vijay Abraham I <kishon@ti.com>
Date:   Thu Jun 3 20:04:26 2021 +0530

    arm64: dts: ti: k3-j721e-common-proc-board: Use external clock for SERDES

    [ Upstream commit f2a7657ad7a821de9cc77d071a5587b243144cd5 ]

    Use external clock for all the SERDES used by PCIe controller. This will
    make the same clock used by the local SERDES as well as the clock
    provided to the PCIe connector.

    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210603143427.28735-4-kishon@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2920f3582a2b382765ed779e5671ddbea877cd17
Author: Kishon Vijay Abraham I <kishon@ti.com>
Date:   Thu Jun 3 20:04:24 2021 +0530

    arm64: dts: ti: k3-j721e-main: Fix external refclk input to SERDES

    [ Upstream commit 5c6d0b55b46aeb91355e6a9616decf50a3778c91 ]

    Rename the external refclk inputs to the SERDES from
    dummy_cmn_refclk/dummy_cmn_refclk1 to cmn_refclk/cmn_refclk1
    respectively. Also move the external refclk DT nodes outside the
    cbass_main DT node. Since in j721e common processor board, only the
    cmn_refclk1 is connected to 100MHz clock, fix the clock frequency.

    Fixes: afd094ebe69f ("arm64: dts: ti: k3-j721e-main: Add WIZ and SERDES PHY nodes")
    Signed-off-by: Kishon Vijay Abraham I <kishon@ti.com>
    Reviewed-by: Aswath Govindraju <a-govindraju@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210603143427.28735-2-kishon@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1ff512116e9bdf45b687ec87b5a2c5750934db5
Author: Cristian Marussi <cristian.marussi@arm.com>
Date:   Tue Jun 8 11:30:56 2021 +0100
```

```
    firmware: arm_scmi: Add delayed response status check

    [ Upstream commit f1748b1ee1fa0fd1a074504045b530b62f949188 ]

    A successfully received delayed response could anyway report a failure at
    the protocol layer in the message status field.

    Add a check also for this error condition.

    Link: https://lore.kernel.org/r/20210608103056.3388-1-cristian.marussi@arm.com
    Fixes: 58ecdf03dbb9 ("firmware: arm_scmi: Add support for asynchronous commands and
delayed response")
    Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
    Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b28d0838e94bac6e6c9e4db799dcf0c3652eaec
Author: Stefan Wahren <stefan.wahren@i2se.com>
Date:   Sat May 29 15:02:51 2021 +0200

    Revert "ARM: dts: bcm283x: increase dwc2's RX FIFO size"

    [ Upstream commit 77daceabedb42482bb6200fa26047c5591716e45 ]

    This reverts commit 278407a53c3b33fb820332c4d39eb39316c3879a.

    The original change breaks USB config on Raspberry Pi Zero and Pi 4 B,
    because it exceeds the total fifo size of 4080. A naive attempt to reduce
    g-tx-fifo-size doesn't help on Raspberry Pi Zero. So better go back.

    Fixes: 278407a53c3b ("ARM: dts: bcm283x: increase dwc2's RX FIFO size")
    Signed-off-by: Stefan Wahren <stefan.wahren@i2se.com>
    Cc: Pavel Hofman <pavel.hofman@ivitera.com>
    Link: https://lore.kernel.org/r/1622293371-5997-1-git-send-email-stefan.wahren@i2se.com
    Signed-off-by: Nicolas Saenz Julienne <nsaenz@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5bb263488bd034fc35951eba82c8284cb3429561
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Tue Jun 1 17:12:50 2021 +0200

    arm64: dts: renesas: r8a779a0: Drop power-domains property from GIC node

    [ Upstream commit 1771a33b34421050c7b830f0a8af703178ba9d36 ]

    "make dtbs_check":

        arm64/boot/dts/renesas/r8a779a0-falcon.dt.yaml: interrupt-controller@f1000000:
'power-domains' does not match any of the regexes: '^(msi-controller|gic-its|interrupt-
controller)@[0-9a-f]+$', '^gic-its@', '^interrupt-controller@[0-9a-f]+$', 'pinctrl-[0-9]+'
                From schema: Documentation/devicetree/bindings/interrupt-controller/arm,gic-
v3.yaml

    Remove the "power-domains" property, as the GIC on R-Car V3U is
    always-on, and not part of a clock domain.

    Fixes: 834c310f541839b6 ("arm64: dts: renesas: Add Renesas R8A779A0 SoC support")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Yoshihiro Shimoda <yoshihiro.shimoda.uh@renesas.com>
    Link:
https://lore.kernel.org/r/a9ae5cbc7c586bf2c6b18ddc665ad7051bd1d206.1622560236.git.geert+rene
sas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit cb5259dacdbc72b63a2a2d282949c1273d648b99
Author: Philipp Zabel <p.zabel@pengutronix.de>
Date:   Mon Jun 7 10:26:15 2021 +0200

    reset: bail if try_module_get() fails

    [ Upstream commit 4fb26fb83f0def3d39c14e268bcd4003aae8fade ]

    Abort instead of returning a new reset control for a reset controller
    device that is going to have its module unloaded.

    Reported-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Fixes: 61fc41317666 ("reset: Add reset controller API")
    Acked-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Link: https://lore.kernel.org/r/20210607082615.15160-1-p.zabel@pengutronix.de
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 39119c5e50a95a4591abf253a3a8b5255cd35ca3
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Wed May 12 15:07:09 2021 +0200

    ARM: dts: BCM5301X: Fixup SPI binding

    [ Upstream commit d5aede3e6dd1b8ca574600a1ecafe1e580c53f2f ]

    1. Reorder interrupts
    2. Fix typo: s/spi_lr_overhead/spi_lr_overread/
    3. Rename node: s/spi-nor@0/flash@0/

    This fixes:
    arch/arm/boot/dts/bcm4709-buffalo-wxr-1900dhp.dt.yaml: spi@18029200: interrupt-names:
'oneOf' conditional failed, one must be fixed:
            ['spi_lr_fullness_reached', 'spi_lr_session_aborted', 'spi_lr_impatient',
'spi_lr_session_done', 'spi_lr_overhead', 'mspi_done', 'mspi_halted'] is too long
            Additional items are not allowed ('spi_lr_session_aborted', 'spi_lr_impatient',
'spi_lr_session_done', 'spi_lr_overhead', 'mspi_done', 'mspi_halted' were unexpected)
            'mspi_done' was expected
            'spi_l1_intr' was expected
            'mspi_halted' was expected
            'spi_lr_fullness_reached' was expected
            'spi_lr_session_aborted' was expected
            'spi_lr_impatient' was expected
            'spi_lr_session_done' was expected
            'spi_lr_overread' was expected
            From schema: Documentation/devicetree/bindings/spi/brcm,spi-bcm-qspi.yaml
    arch/arm/boot/dts/bcm4709-buffalo-wxr-1900dhp.dt.yaml: spi-nor@0: $nodename:0: 'spi-
nor@0' does not match '^flash(@.*)?$'
            From schema: Documentation/devicetree/bindings/mtd/jedec,spi-nor.yaml

    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2c62ef35e795c83c36806490d2c44c893b03690
Author: Nicolas Ferre <nicolas.ferre@microchip.com>
Date:   Fri Jun 4 15:19:25 2021 +0200

    dt-bindings: i2c: at91: fix example for scl-gpios

    [ Upstream commit 92e669017ff1616ba7d8ba3c65f5193bc2a7acbe ]

    The SCL gpio pin used by I2C bus for recovery needs to be configured as
    open drain, so fix the binding example accordingly.
    In relation with fix c5a283802573 ("ARM: dts: at91: Configure I2C SCL
```

```
      gpio as open drain").

      Signed-off-by: Nicolas Ferre <nicolas.ferre@microchip.com>
      Fixes: 19e5cef058a0 ("dt-bindings: i2c: at91: document optional bus recovery
  properties")
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 365987eb0d0ee3aed11424729e02857d7d7fd51e
  Author: Cristian Marussi <cristian.marussi@arm.com>
  Date:   Tue Jun 1 11:24:17 2021 +0100

      firmware: arm_scmi: Reset Rx buffer to max size during async commands

      [ Upstream commit 0cb7af474e0dbb2f500c67aa62b6db9fafa74de2 ]

      During an async commands execution the Rx buffer length is at first set
      to max_msg_sz when the synchronous part of the command is first sent.
      However once the synchronous part completes the transport layer waits
      for the delayed response which will be processed using the same xfer
      descriptor initially allocated. Since synchronous response received at
      the end of the xfer will shrink the Rx buffer length to the effective
      payload response length, it needs to be reset again.

      Raise the Rx buffer length again to max_msg_sz before fetching the
      delayed response to ensure full response is read correctly from the
      shared memory.

      Link: https://lore.kernel.org/r/20210601102421.26581-2-cristian.marussi@arm.com
      Fixes: 58ecdf03dbb9 ("firmware: arm_scmi: Add support for asynchronous commands and
  delayed response")
      Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
      [sudeep.holla: moved reset to scmi_handle_response as it could race with
                     do_xfer_with_response]
      Signed-off-by: Sudeep Holla <sudeep.holla@arm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 1a35a7db34304bde8560d49d2ad6011c0152fb97
  Author: Weiyi Lu <weiyi.lu@mediatek.com>
  Date:   Tue Jun 1 11:59:03 2021 +0800

      soc: mtk-pm-domains: Fix the clock prepared issue

      [ Upstream commit f0fce06e345dc4f75c1cdd21840780f5fe2df1f3 ]

      In this new power domain driver, when adding one power domain
      it will prepare the dependent clocks at the same.
      So we only do clk_bulk_enable/disable control during power ON/OFF.
      When system suspend, the pm runtime framework will forcely power off
      power domains. However, the dependent clocks are disabled but kept
      prepared.

      In MediaTek clock drivers, PLL would be turned ON when we do
      clk_bulk_prepare control.

      Clock hierarchy:
      PLL -->
             DIV_CK -->
                       CLK_MUX
                       (may be dependent clocks)
                              -->
                                  SUBSYS_CG
                                  (may be dependent clocks)

      It will lead some unexpected clock states during system suspend.
      This patch will fix by doing prepare_enable/disable_unprepare on
```

```
    dependent clocks at the same time while we are going to power on/off
    any power domain.

    Fixes: 59b644b01cf4 ("soc: mediatek: Add MediaTek SCPSYS power domains")
    Signed-off-by: Weiyi Lu <weiyi.lu@mediatek.com>
    Signed-off-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Reviewed-by: chun-jie.chen <chun-jie.chen@mediatek.com>
    Reviewed-by: Enric Balletbo i Serra <enric.balletbo@collabora.com>
    Link: https://lore.kernel.org/r/20210601035905.2970384-1-hsinyi@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 02fb5809342a42953ebe7ea3e909d45990067feb
Author: Hsin-Yi Wang <hsinyi@chromium.org>
Date:   Tue Jun 1 11:59:04 2021 +0800

    soc: mtk-pm-domains: do not register smi node as syscon

    [ Upstream commit eed6ff1bb2da65067d928f4ab322c7d75f944fa4 ]

    Mediatek requires mmsys clocks to be unprepared during suspend,
    otherwise system has chances to hang.

    syscon_regmap_lookup_by_phandle_optional() will attach and prepare the
    first clock in smi node, leading to additional prepare to the clock
    which is not balanced with the prepare/unprepare pair in resume/suspend
    callbacks.

    If a power domain node requests an smi node and the smi node's first
    clock is an mmsys clock, it will results in an unstable suspend resume.

    Fixes: f414854c8843 ("soc: mediatek: pm-domains: Add SMI block as bus protection block")
    Signed-off-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Reviewed-by: chun-jie.chen <chun-jie.chen@mediatek.com>
    Reviewed-by: Enric Balletbo i Serra <enric.balletbo@collabora.com>
    Link: https://lore.kernel.org/r/20210601035905.2970384-2-hsinyi@chromium.org
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ea6ad20a538157de59dfda42383693e6ba945b12
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Thu May 13 21:26:46 2021 +0800

    firmware: tegra: Fix error return code in tegra210_bpmp_init()

    [ Upstream commit 7fea67710e9f6a111a2c9440576f2396ccd92d57 ]

    When call irq_get_irq_data() to get the IRQ's irq_data failed, an
    appropriate error code -ENOENT should be returned. However, we directly
    return 'err', which records the IRQ number instead of the error code.

    Fixes: 139251fc2208 ("firmware: tegra: add bpmp driver for Tegra210")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ef47def7bbaae0ef2b04b228abb9a3cb51a2811
Author: Douglas Anderson <dianders@chromium.org>
Date:   Mon Mar 15 10:38:54 2021 -0700

    arm64: dts: qcom: sc7180: Fix sc7180-qmp-usb3-dp-phy reg sizes

    [ Upstream commit c1124180eb9883891ad2acef89c9d17d6190eab4 ]
```

```
    As per Dmitry Baryshkov [1]:
    a) The 2nd "reg" should be 0x3c because "Offset 0x38 is
       USB3_DP_COM_REVISION_ID3 (not used by the current driver though)."
    b) The 3rd "reg" "is a serdes region and qmp_v3_dp_serdes_tbl contains
       registers 0x148 and 0x154."

    I think because the 3rd "reg" is a serdes region we should just use
    the same size as the 1st "reg"?

    [1] https://lore.kernel.org/r/ee5695bb-a603-0dd5-7a7f-695e919b1af1@linaro.org

    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Stephen Boyd <swboyd@chromium.org>
    Cc: Stephen Boyd <swboyd@chromium.org>
    Cc: Jeykumar Sankaran <jsanka@codeaurora.org>
    Cc: Chandan Uddaraju <chandanu@codeaurora.org>
    Cc: Vara Reddy <varar@codeaurora.org>
    Cc: Tanmay Shah <tanmay@codeaurora.org>
    Cc: Rob Clark <robdclark@chromium.org>
    Fixes: 58fd7ae621e7 ("arm64: dts: qcom: sc7180: Update dts for DP phy inside QMP phy")
    Reported-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Link:
https://lore.kernel.org/r/20210315103836.1.I9a97120319d43b42353aeac4d348624d60687df7@changei
d
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 77cb5b1b127f9709cd59f7eb8565efa2e8bd568c
Author: Stephen Boyd <swboyd@chromium.org>
Date:   Wed Mar 24 16:14:24 2021 -0700

    arm64: dts: qcom: c630: Add no-hpd to DSI bridge node

    [ Upstream commit c0dcfe6a784fdf7fcc0fdc74bfbb06e9f77de964 ]

    We should indicate that we're not using the HPD pin on this device, per
    the binding document. Otherwise if code in the future wants to enable
    HPD in the bridge when this property is absent we'll be enabling HPD
    when it isn't supposed to be used. Presumably this board isn't using hpd
    on the bridge.

    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Cc: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Cc: Douglas Anderson <dianders@chromium.org>
    Cc: Steev Klimaszewski <steev@kali.org>
    Fixes: 956e9c85f47b ("arm64: dts: qcom: c630: Define eDP bridge and panel")
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Link: https://lore.kernel.org/r/20210324231424.2890039-1-swboyd@chromium.org
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c8072f1620ce0025914653db6f70372d7428866a
Author: Stephen Boyd <swboyd@chromium.org>
Date:   Tue Mar 23 19:55:34 2021 -0700

    arm64: dts: qcom: trogdor: Add no-hpd to DSI bridge node

    [ Upstream commit 5f551b5ce55575b14c26933fe9b49365ea246b3d ]

    We should indicate that we're not using the HPD pin on this device, per
    the binding document. Otherwise if code in the future wants to enable
    HPD in the bridge when this property is absent we'll be wasting power
    powering hpd when we don't use it on trogdor boards. We didn't notice
    this before because the kernel driver blindly disables hpd, but that
```

```
    won't be true for much longer.

    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Cc: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Cc: Douglas Anderson <dianders@chromium.org>
    Fixes: 7ec3e67307f8 ("arm64: dts: qcom: sc7180-trogdor: add initial trogdor and lazor
dt")
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Link: https://lore.kernel.org/r/20210324025534.1837405-1-swboyd@chromium.org
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d2b0b26399f3dc04118bf116edccf661e206a9e
Author: Marek Vasut <marex@denx.de>
Date:   Fri Apr 9 01:00:01 2021 +0200

    ARM: dts: stm32: Rework LAN8710Ai PHY reset on DHCOM SoM

    [ Upstream commit 1cebcf9932ab76102e8cfc555879574693ba8956 ]

    The Microchip LAN8710Ai PHY requires XTAL1/CLKIN external clock to be
    enabled when the nRST is toggled according to datasheet Microchip
    LAN8710A/LAN8710Ai DS00002164B page 35 section 3.8.5.1 Hardware Reset:
      "
      A Hardware reset is asserted by driving the nRST input pin low. When
      driven, nRST should be held low for the minimum time detailed in
      Section 5.5.3, "Power-On nRST & Configuration Strap Timing," on page
      59 to ensure a proper transceiver reset. During a Hardware reset, an
      external clock must be supplied to the XTAL1/CLKIN signal.
      "
    This is accidentally fulfilled in the current setup, where ETHCK_K is used
    to supply both PHY XTAL1/CLKIN and is also fed back through eth_clk_fb to
    supply ETHRX clock of the DWMAC. Hence, the DWMAC enables ETHRX clock,
    that has ETHCK_K as parent, so ETHCK_K clock are also enabled, and then
    the PHY reset toggles.

    However, this is not always the case, e.g. in case the PHY XTAL1/CLKIN
    clock are supplied by some other clock source than ETHCK_K or in case
    ETHRX clock are not supplied by ETHCK_K. In the later case, ETHCK_K would
    be kept disabled, while ETHRX clock would be enabled, so the PHY would
    not be receiving XTAL1/CLKIN clock and the reset would fail.

    Improve the DT by adding the PHY clock phandle into the PHY node, which
    then also requires moving the PHY reset GPIO specifier in the same place
    and that then also requires correct PHY reset GPIO timing, so add that
    too.

    A brief note regarding the timing, the datasheet says the reset should
    stay asserted for at least 100uS and software should wait at least 200nS
    after deassertion. Set both delays to 500uS which should be plenty.

    Fixes: 34e0c7847dcf ("ARM: dts: stm32: Add DH Electronics DHCOM STM32MP1 SoM and PDK2
board")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Cc: Alexandre Torgue <alexandre.torgue@st.com>
    Cc: Patrice Chotard <patrice.chotard@st.com>
    Cc: Patrick Delaunay <patrick.delaunay@st.com>
    Cc: linux-stm32@st-md-mailman.stormreply.com
    To: linux-arm-kernel@lists.infradead.org
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3d54e663d3b6347e34c3b2892266214613cbb67
Author: Geert Uytterhoeven <geert+renesas@glider.be>
```

```
Date:   Thu Apr 29 14:39:12 2021 +0200

    arm64: dts: renesas: r8a7796[01]: Fix OPP table entry voltages

    [ Upstream commit 659b38203f04f5c3d1dc60f1a3e54b582ad3841c ]

    Correct the voltages in the "Power Optimized" (<= 1.5 GHz) Cortex-A57
    operating point table entries for the R-Car M3-W and M3-W+ SoCs from
    0.82V to 0.83V, as per the R-Car Gen3 EC Manual Errata for Revision
    0.53.

    Based on a patch for R-Car M3-W in the BSP by Takeshi Kihara
    <takeshi.kihara.df@renesas.com>.

    Fixes: da7e3113344fda50 ("arm64: dts: renesas: r8a7796: Add OPPs table for cpu devices")
    Fixes: f51746ad7d1ff6b4 ("arm64: dts: renesas: Add Renesas R8A77961 SoC support")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/r/b9e9db907514790574429b83d070c823b36085ef.1619699909.git.geert+rene
sas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8303239e753a0c99e26f7098f4acc9f4513c79e
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Thu May 27 15:42:42 2021 +0200

    arm64: dts: renesas: Add missing opp-suspend properties

    [ Upstream commit 44b615ac9fab16d1552cd8360454077d411e3c35 ]

    Tag the highest "Power Optimized" (1.5 GHz) Cortex-A57 operating point
    table entries for the RZ/G2M, R-Car M3-W and M3-W+ SoCs with the
    "opp-suspend" property.  This makes sure the system will enter suspend
    in the same performance state as it will be resumed by the firmware
    later, avoiding state inconsistencies after resume.

    Based on a patch for R-Car M3-W in the BSP by Takeshi Kihara
    <takeshi.kihara.df@renesas.com>.

    Fixes: 800037e815b91d8c ("arm64: dts: renesas: r8a774a1: Add operating points")
    Fixes: da7e3113344fda50 ("arm64: dts: renesas: r8a7796: Add OPPs table for cpu devices")
    Fixes: f51746ad7d1ff6b4 ("arm64: dts: renesas: Add Renesas R8A77961 SoC support")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Niklas Söderlund <niklas.soderlund+renesas@ragnatech.se>
    Link:
https://lore.kernel.org/r/45a061c3b0463aac7d10664f47c4afdd999da50d.1619699721.git.geert+rene
sas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bd8eefc0b53d062a79fb38806c3fe24c5483b7f5
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed May 12 10:31:41 2021 +0530

    ARM: dts: qcom: sdx55-telit: Represent secure-regions as 64-bit elements

    [ Upstream commit 0fa1baeedf06765ec6b441692ba2a2e83b7d17dc ]

    The corresponding MTD code expects the regions to be of 64-bit elements.
    Hence, prefix "/bits/ 64", otherwise the regions will not be parsed
    correctly.

    Fixes: 6a5d3c611930 ("ARM: dts: qcom: sdx55: Add basic devicetree support for Telit
FN980 TLB")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20210512050141.43338-2-manivannan.sadhasivam@linaro.org
```

```
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6dfc3b366e20b9e5438f4b0ea0e0ec0c37a7f2a3
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Wed May 12 10:31:40 2021 +0530

    ARM: dts: qcom: sdx55-t55: Represent secure-regions as 64-bit elements

    [ Upstream commit 619d3c4bf8f346ac9192d3c266efc9e231ca5d17 ]

    The corresponding MTD code expects the regions to be of 64-bit elements.
    Hence, prefix "/bits/ 64", otherwise the regions will not be parsed
    correctly.

    Fixes: 3263d4be5788 ("ARM: dts: qcom: sdx55: Add basic devicetree support for
Thundercomm T55")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20210512050141.43338-1-manivannan.sadhasivam@linaro.org
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b5c14a6bbe4537f5deaf6d05be30b5f609c5432
Author: Roger Quadros <rogerq@ti.com>
Date:   Wed May 12 21:03:08 2021 +0530

    arm64: dts: ti: j7200-main: Enable USB2 PHY RX sensitivity workaround

    [ Upstream commit a2894d85f44ba3f2bdf5806c8dc62e2ec40c1c09 ]

    Enable work around feature built into the controller to address issue with
    RX Sensitivity for USB2 PHY.

    Fixes: 6197d7139d12 ("arm64: dts: ti: k3-j7200-main: Add USB controller")
    Signed-off-by: Roger Quadros <rogerq@ti.com>
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Reviewed-by: Vignesh Raghavendra <vigneshr@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210512153308.5840-1-a-govindraju@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ec80ffdc4283b7aa75936a4f8b486a5b8852ca65
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Fri Apr 23 12:17:57 2021 +0530

    arm64: dts: ti: k3-j7200: Remove "#address-cells" property from GPIO DT nodes

    [ Upstream commit 6ec8ba764165f6ecb6f6f7efbfd2ec7ad76dedcb ]

    GPIO device tree nodes do not have child nodes. Therefore, "#address-cells"
    property should not be added.

    Fixes: e0b2e6af39ea ("arm64: dts: ti: k3-j7200: Add gpio nodes")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Reviewed-by: Lokesh  Vutla <lokeshvutla@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210423064758.25520-1-a-govindraju@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d300d33d5674c54a7bc036c29bbda07fee9c864
Author: Aswath Govindraju <a-govindraju@ti.com>
Date:   Fri Apr 23 11:31:33 2021 +0530

    arm64: dts: ti: k3-am64-mcu: Fix the compatible string in GPIO DT node
```

```
    [ Upstream commit ec2fb989d03e7f79f7cd901cf9abf40aebba7acf ]

    Fix the compatible string in mcu domain GPIO device tree node.

    Fixes: 01a91e01b8fd ("arm64: dts: ti: k3-am64: Add GPIO DT nodes")
    Signed-off-by: Aswath Govindraju <a-govindraju@ti.com>
    Reviewed-by: Lokesh Vutla <lokeshvutla@ti.com>
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Link: https://lore.kernel.org/r/20210423060133.16473-1-a-govindraju@ti.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1b67a18629234664f3afd65bb3559fb8a3c49b8
Author: Caleb Connolly <caleb@connolly.tech>
Date:   Sun May 2 01:42:57 2021 +0000

    arm64: dts: qcom: sdm845-oneplus-common: guard rmtfs-mem

    [ Upstream commit e60fd5ac1f6851be5b2c042b39584bfcf8a66f57 ]

    The rmtfs_mem region is a weird one, downstream allocates it
    dynamically, and supports a "qcom,guard-memory" property which when set
    will reserve 4k above and below the rmtfs memory.

    A common from qcom 4.9 kernel msm_sharedmem driver:

    /*
     * If guard_memory is set, then the shared memory region
     * will be guarded by SZ_4K at the start and at the end.
     * This is needed to overcome the XPU limitation on few
     * MSM HW, so as to make this memory not contiguous with
     * other allocations that may possibly happen from other
     * clients in the system.
    */

    When the kernel tries to touch memory that is too close the
    rmtfs region it may cause an XPU violation. Such is the case on the
    OnePlus 6 where random crashes would occur usually after boot.

    Reserve 4k above and below the rmtfs_mem to avoid hitting these XPU
    Violations.

    This doesn't entirely solve the random crashes on the OnePlus 6/6T but
    it does seem to prevent the ones which happen shortly after modem
    bringup.

    Fixes: 288ef8a42612 ("arm64: dts: sdm845: add oneplus6/6t devices")
    Signed-off-by: Caleb Connolly <caleb@connolly.tech>
    Link: https://lore.kernel.org/r/20210502014146.85642-4-caleb@connolly.tech
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d8ef498abc672bcbb30f13e52292ac517e6a4cad
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Thu Apr 29 14:41:15 2021 +0200

    ARM: dts: r8a7779, marzen: Fix DU clock names

    [ Upstream commit 6ab8c23096a29b69044209a5925758a6f88bd450 ]

    "make dtbs_check" complains:

        arch/arm/boot/dts/r8a7779-marzen.dt.yaml: display@fff80000: clock-names:0: 'du.0'
was expected

    Change the first clock name to match the DT bindings.
```

```
    This has no effect on actual operation, as the Display Unit driver in
    Linux does not use the first clock name on R-Car H1, but just grabs the
    first clock.

    Fixes: 665d79aa47cb3983 ("ARM: shmobile: marzen: Add DU external pixel clock to DT")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Link:
https://lore.kernel.org/r/9d5e1b371121883b3b3e10a3df43802a29c6a9da.1619699965.git.geert+rene
sas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d32c0007098a1e1cf603fc9c953c6bce42c1231
Author: Valentine Barshak <valentine.barshak@cogentembedded.com>
Date:   Fri Mar 26 13:10:50 2021 +0100

    arm64: dts: renesas: v3msk: Fix memory size

    [ Upstream commit a422ec20caef6a50cf3c1efa93538888ebd576a6 ]

    The V3MSK board has 2 GiB RAM according to the datasheet and schematics.

    Signed-off-by: Valentine Barshak <valentine.barshak@cogentembedded.com>
    [geert: Verified schematics]
    Fixes: cc3e267e9bb0ce7f ("arm64: dts: renesas: initial V3MSK board device tree")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20210326121050.1578460-1-geert+renesas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 104f7241f12de24f17ccd5aedbfd56a133c78bb3
Author: Dan Carpenter <dan.carpenter@oracle.com>
Date:   Tue May 11 10:19:26 2021 +0300

    rtc: fix snprintf() checking in is_rtc_hctosys()

    [ Upstream commit 54b909436ede47e0ee07f1765da27ec2efa41e84 ]

    The scnprintf() function silently truncates the printf() and returns
    the number bytes that it was able to copy (not counting the NUL
    terminator).  Thus, the highest value it can return here is
    "NAME_SIZE - 1" and the overflow check is dead code.  Fix this by
    using the snprintf() function which returns the number of bytes that
    would have been copied if there was enough space and changing the
    condition from "> NAME_SIZE" to ">= NAME_SIZE".

    Fixes: 92589c986b33 ("rtc-proc: permit the /proc/driver/rtc device to use other
devices")
    Signed-off-by: Dan Carpenter <dan.carpenter@oracle.com>
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Link: https://lore.kernel.org/r/YJov/pcGmhLi2pEl@mwanda
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b46343d52e86c5d5f5f23698e7faf85146a7856
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Thu Apr 22 17:15:21 2021 +0200

    rtc: bd70528: fix BD71815 watchdog dependency

    [ Upstream commit b0ddc5b170058a9ed3c9f031501d735a4eb8ee89 ]

    The added Kconfig dependency is slightly incorrect, which can
    lead to a link failure when the watchdog is a loadable module:

    arm-linux-gnueabi-ld: drivers/rtc/rtc-bd70528.o: in function
`bd70528_set_rtc_based_timers':
```

```
      rtc-bd70528.c:(.text+0x6cc): undefined reference to `bd70528_wdt_set'
      arm-linux-gnueabi-ld: drivers/rtc/rtc-bd70528.o: in function `bd70528_set_time':
      rtc-bd70528.c:(.text+0xaa0): undefined reference to `bd70528_wdt_lock'
      arm-linux-gnueabi-ld: rtc-bd70528.c:(.text+0xab8): undefined reference to
  `bd70528_wdt_unlock'
      arm-linux-gnueabi-ld: drivers/rtc/rtc-bd70528.o: in function `bd70528_alm_enable':
      rtc-bd70528.c:(.text+0xfc0): undefined reference to `bd70528_wdt_lock'
      arm-linux-gnueabi-ld: rtc-bd70528.c:(.text+0x1030): undefined reference to
  `bd70528_wdt_unlock'

      The problem is that it allows to be built-in if MFD_ROHM_BD71828
      is built-in, even when the watchdog is a loadable module.

      Rework this so that having the watchdog as a loadable module always
      forces the rtc to be a module as well instead of built-in,
      regardless of bd71828.

      Fixes: c56dc069f268 ("rtc: bd70528: Support RTC on ROHM BD71815")
      Signed-off-by: Arnd Bergmann <arnd@arndb.de>
      Reviewed-by: Guenter Roeck <linux@roeck-us.net>
      Reviewed-by: Matti Vaittinen <matti.vaittinen@fi.rohmeurope.com>
      Acked-by: Randy Dunlap <rdunlap@infradead.org> # build-tested
      Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
      Link: https://lore.kernel.org/r/20210422151545.2403356-1-arnd@kernel.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit dff5eb6da4f8a13fcca79d1129f12dd09dd75daf
  Author: Salvatore Bonaccorso <carnil@debian.org>
  Date:   Mon May 24 14:21:11 2021 +0200

      ARM: dts: sun8i: h3: orangepi-plus: Fix ethernet phy-mode

      [ Upstream commit b19d3479f25e8a0ff24df0b46c82e50ef0f900dd ]

      Commit bbc4d71d6354 ("net: phy: realtek: fix rtl8211e rx/tx delay
      config") sets the RX/TX delay according to the phy-mode property in the
      device tree. For the Orange Pi Plus board this is "rgmii", which is the
      wrong setting.

      Following the example of a900cac3750b ("ARM: dts: sun7i: a20: bananapro:
      Fix ethernet phy-mode") the phy-mode is changed to "rgmii-id" which gets
      the Ethernet working again on this board.

      Fixes: bbc4d71d6354 ("net: phy: realtek: fix rtl8211e rx/tx delay config")
      Reported-by: "B.R. Oake" <broake@mailfence.com>
      Reported-by: Vagrant Cascadian <vagrant@reproducible-builds.org>
      Link: https://bugs.debian.org/988574
      Signed-off-by: Salvatore Bonaccorso <carnil@debian.org>
      Signed-off-by: Maxime Ripard <maxime@cerno.tech>
      Link: https://lore.kernel.org/r/20210524122111.416885-1-carnil@debian.org
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 7540e4352365a5cd9371e812c9540a690913a862
  Author: Zhen Lei <thunder.leizhen@huawei.com>
  Date:   Sat May 15 12:00:04 2021 +0800

      memory: pl353: Fix error return code in pl353_smc_probe()

      [ Upstream commit 76e5624f3f9343a621dd3f4006f4e4d9c3f91e33 ]

      When no child nodes are matched, an appropriate error code -ENODEV should
      be returned. However, we currently do not explicitly assign this error
      code to 'err'. As a result, 0 was incorrectly returned.

      Fixes: fee10bd22678 ("memory: pl353: Add driver for arm pl353 static memory controller")
```

```
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Link: https://lore.kernel.org/r/20210515040004.6983-1-thunder.leizhen@huawei.com
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 410c26a81b6c05c5bf0e0f359bc3c6ae6d875213
Author: Enric Balletbo i Serra <enric.balletbo@collabora.com>
Date:   Fri Apr 23 09:52:01 2021 +0200

    arm64: defconfig: Do not override the MTK_PMIC_WRAP symbol

    [ Upstream commit f0e70d4946332c681ceaba940652f30c7c33473d ]

    Commit 'fbbe38309d56 ("arm64: defconfig: Allow mt8173-based boards to boot
    from usb")' added the MTK_PMIC_WRAP config built-in. It needs to be
    built-in in order to be able to boot from USB or the MMC without needing
    a ramdisk, but that symbol was already defined as a module so now we are
    getting the following warning:

      arch/arm64/configs/defconfig:996:warning: override: reassigning to symbol
MTK_PMIC_WRAP

    Remove the MTK_PMIC_WRAP=m from the defconfig to remove the error.

    Fixes: fbbe38309d56 ("arm64: defconfig: Allow mt8173-based boards to boot from usb")
    Signed-off-by: Enric Balletbo i Serra <enric.balletbo@collabora.com>
    Link: https://lore.kernel.org/r/20210423075201.2616023-1-enric.balletbo@collabora.com
    Signed-off-by: Matthias Brugger <matthias.bgg@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 91253471596e7faf14dc6efa7f3156e771d80b9b
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 11:14:43 2021 +0800

    reset: brcmstb: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit e207457f9045343a24d936fbb67eb4b412f1c6ad ]

    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>
    Fixes: 77750bc089e4 ("reset: Add Broadcom STB SW_INIT reset controller driver")
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Link: https://lore.kernel.org/r/1620789283-15048-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19c0f8293f31ba223499b8922c49ceaafa18f76f
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Fri Apr 23 12:18:15 2021 +0200

    memory: atmel-ebi: add missing of_node_put for loop iteration

    [ Upstream commit 907c5bbb514a4676160e79764522fff56ce3448e ]

    Early exits from for_each_available_child_of_node() should decrement the
    node reference counter.  Reported by Coccinelle:

      drivers/memory/atmel-ebi.c:593:1-33: WARNING:
        Function "for_each_available_child_of_node" should have of_node_put() before return
around line 604.
```

```
    Fixes: 6a4ec4cd0888 ("memory: add Atmel EBI (External Bus Interface) driver")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210423101815.119341-2-
krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e075ad208d19a1a671a8830f1fb6096268044be6
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Fri Apr 23 12:18:14 2021 +0200

    memory: stm32-fmc2-ebi: add missing of_node_put for loop iteration

    [ Upstream commit 2f9dc6a357ff3b82c1e54d29fb5d52b8d4a0c587 ]

    Early exits from for_each_available_child_of_node() should decrement the
    node reference counter.  Reported by Coccinelle:

      drivers/memory/stm32-fmc2-ebi.c:1046:1-33: WARNING:
        Function "for_each_available_child_of_node" should have of_node_put() before return
around line 1051.

    Fixes: 66b8173a197f ("memory: stm32-fmc2-ebi: add STM32 FMC2 EBI controller driver")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Reviewed-by: Christophe Kerello <christophe.kerello@foss.st.com>
    Link: https://lore.kernel.org/r/20210423101815.119341-1-
krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 444c3049db1073f588b1646af6e1fab1bacc0f74
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Wed May 5 09:59:41 2021 -0400

    ARM: dts: exynos: fix PWM LED max brightness on Odroid XU4

    [ Upstream commit fd2f1717966535b7d0b6fe45cf0d79e94330da5f ]

    There is no "max_brightness" property as pointed out by dtschema:

      arch/arm/boot/dts/exynos5422-odroidxu4.dt.yaml: led-controller: led-1: 'max-
brightness' is a required property

    Fixes: 6658356014cb ("ARM: dts: Add support Odroid XU4 board for exynos5422-odroidxu4")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210505135941.59898-5-krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 209aa4be8b40698cbce5b0a41ae3a3eb4f104fed
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Wed May 5 09:59:40 2021 -0400

    ARM: dts: exynos: fix PWM LED max brightness on Odroid HC1

    [ Upstream commit a7e59c84cf2055a1894f45855c8319191f2fa59e ]

    There is no "max_brightness" property as pointed out by dtschema:

      arch/arm/boot/dts/exynos5422-odroidhc1.dt.yaml: led-controller: led-1: 'max-
brightness' is a required property

    Fixes: 1ac49427b566 ("ARM: dts: exynos: Add support for Hardkernel's Odroid HC1 board")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210505135941.59898-4-krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bbfd45abcd39b038a45ba2f2ae9ba9b5f9d1a1e
```

```
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Wed May 5 09:59:39 2021 -0400

    ARM: dts: exynos: fix PWM LED max brightness on Odroid XU/XU3

    [ Upstream commit 75121e1dc9fe4def41e63d57f6a53749b88006ed ]

    There is no "max_brightness" property.  This brings the intentional
    brightness reduce of green LED and dtschema checks as well:

      arch/arm/boot/dts/exynos5410-odroidxu.dt.yaml: led-controller-1: led-1: 'max-
brightness' is a required property

    Fixes: 719f39fec586 ("ARM: dts: exynos5422-odroidxu3: Hook up PWM and use it for LEDs")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210505135941.59898-3-krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bd497c174c53252eb984d479f83f7cba90b5d34
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Sun Apr 25 19:49:45 2021 +0200

    ARM: exynos: add missing of_node_put for loop iteration

    [ Upstream commit 48d551bf20858240f38a0276be3016ff379918ac ]

    Early exits from for_each_compatible_node() should decrement the
    node reference counter.  Reported by Coccinelle:

      arch/arm/mach-exynos/exynos.c:52:1-25: WARNING:
        Function "for_each_compatible_node" should have of_node_put() before break around
line 58.

    Fixes: b3205dea8fbf ("ARM: EXYNOS: Map SYSRAM through generic DT bindings")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20210425174945.164612-1-
krzysztof.kozlowski@canonical.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7a91d9c6753190672288537a078689ffc5ea61e9
Author: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
Date:   Fri May 7 07:28:03 2021 -0400

    reset: a10sr: add missing of_match_table reference

    [ Upstream commit 466ba3c8ff4fae39e455ff8d080b3d5503302765 ]

    The driver defined of_device_id table but did not use it with
    of_match_table.  This prevents usual matching via devicetree and causes
    a W=1 warning:

      drivers/reset/reset-a10sr.c:111:34: warning:
        ‘a10sr_reset_of_match’ defined but not used [-Wunused-const-variable=]

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: 627006820268 ("reset: Add Altera Arria10 SR Reset Controller")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@canonical.com>
    Link: https://lore.kernel.org/r/20210507112803.20012-1-krzysztof.kozlowski@canonical.com
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6cf18d4e190b2c51ca555bd6666000c7f96a3b4e
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Wed Mar 31 10:15:19 2021 +0200
```

```
    reset: RESET_INTEL_GW should depend on X86

    [ Upstream commit 6ab9d6219f86f0db916105444813aafce626a2f4 ]

    The Intel Gateway reset controller is only present on Intel Gateway
    platforms.  Hence add a dependency on X86, to prevent asking the user
    about this driver when configuring a kernel without Intel Gateway
    support.

    Fixes: c9aef213e38cde27 ("reset: intel: Add system reset controller driver")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1347d04d66f7bb201a24a7d1b7560174ad8cdefd
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Tue Mar 16 14:37:33 2021 +0100

    reset: RESET_BRCMSTB_RESCAL should depend on ARCH_BRCMSTB

    [ Upstream commit 42f6a76fbe85e5243f83a3ed76809b1ebbb7087e ]

    The Broadcom STB RESCAL reset controller is only present on Broadcom
    BCM7216 platforms.  Hence add a dependency on ARCH_BRCMSTB, to prevent
    asking the user about this driver when configuring a kernel without
    BCM7216 support.

    Also, merely enabling CONFIG_COMPILE_TEST should not enable additional
    code, and thus should not enable this driver by default.

    Fixes: 4cf176e52397853e ("reset: Add Broadcom STB RESCAL reset controller")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Acked-by: Florian Fainelli <f.fainelli@gmail.com>
    Signed-off-by: Philipp Zabel <p.zabel@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b5eb8cd04a2c061a207796cd547a058cbf97cee
Author: Chen-Yu Tsai <wens@csie.org>
Date:   Mon Apr 26 17:59:16 2021 +0800

    arm64: dts: rockchip: Drop fephy pinctrl from gmac2phy on rk3328 rock-pi-e

    [ Upstream commit e6526f90696e6a7d722d04b958f15b97d6fd9ce6 ]

    Turns out the fephy pins are already claimed in the phy node, which is
    rightfully where they should be claimed.

    Drop the pinctrl properties from the gmac2phy node for the ROCK Pi E.

    Fixes: b918e81f2145 ("arm64: dts: rockchip: rk3328: Add Radxa ROCK Pi E")
    Signed-off-by: Chen-Yu Tsai <wens@csie.org>
    Link: https://lore.kernel.org/r/20210426095916.14574-1-wens@kernel.org
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 62ace4c8a6fcdb98f98562579ab7f35fe36cbccb
Author: Tianling Shen <cnsztl@gmail.com>
Date:   Mon Apr 26 19:46:52 2021 +0800

    arm64: dts: rockchip: rename LED label for NanoPi R4S

    [ Upstream commit 6a11ffc2cc54d89719d5b2f3ca44244cebd7ed2e ]

    However "sys" is not a valid function, and it is always on.
```

```
    Let's keep existing functions.

    Fixes: db792e9adbf85f ("rockchip: rk3399: Add support for FriendlyARM NanoPi R4S")

    Suggested-by: Pavel Machek <pavel@ucw.cz>
    Signed-off-by: Tianling Shen <cnsztl@gmail.com>
    Acked-by: Pavel Machek <pavel@ucw.cz>
    Link: https://lore.kernel.org/r/20210426114652.29542-1-cnsztl@gmail.com
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f499d706e21e478a76eed210fdc4cf88c23ae671
Author: Corentin Labbe <clabbe@baylibre.com>
Date:   Wed Apr 28 18:54:57 2021 +0000

    ARM: dts: gemini-rut1xx: remove duplicate ethernet node

    [ Upstream commit 3d3bb3d27cd371d3edb43eeb1beb8ae4e92a356d ]

    Two ethernet node was added by
    commit 95220046a62c ("ARM: dts: Add ethernet to a bunch of platforms")
    and commit d6d0cef55e5b ("ARM: dts: Add the FOTG210 USB host to Gemini boards")

    This patch removes the duplicate one.

    Fixes: d6d0cef55e5b ("ARM: dts: Add the FOTG210 USB host to Gemini boards")
    Signed-off-by: Corentin Labbe <clabbe@baylibre.com>
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 700ecdcb3e057a3e17ba921cda91bb8940232278
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed Jul 7 18:07:41 2021 -0700

    hexagon: use common DISCARDS macro

    [ Upstream commit 681ba73c72302214686401e707e2087ed11a6556 ]

    ld.lld warns that the '.modinfo' section is not currently handled:

    ld.lld: warning: kernel/built-in.a(workqueue.o):(.modinfo) is being placed in '.modinfo'
    ld.lld: warning: kernel/built-in.a(printk/printk.o):(.modinfo) is being placed in
'.modinfo'
    ld.lld: warning: kernel/built-in.a(irq/spurious.o):(.modinfo) is being placed in
'.modinfo'
    ld.lld: warning: kernel/built-in.a(rcu/update.o):(.modinfo) is being placed in
'.modinfo'

    The '.modinfo' section was added in commit 898490c010b5 ("moduleparam:
    Save information about built-in modules in separate file") to the DISCARDS
    macro but Hexagon has never used that macro.  The unification of DISCARDS
    happened in commit 023bf6f1b8bf ("linker script: unify usage of discard
    definition") in 2009, prior to Hexagon being added in 2011.

    Switch Hexagon over to the DISCARDS macro so that anything that is
    expected to be discarded gets discarded.

    Link: https://lkml.kernel.org/r/20210521011239.1332345-3-nathan@kernel.org
    Fixes: e95bf452a9e2 ("Hexagon: Add configuration and makefiles for the Hexagon
architecture.")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Acked-by: Brian Cain <bcain@codeaurora.org>
    Cc: David Rientjes <rientjes@google.com>
    Cc: Oliver Glitta <glittao@gmail.com>
```

```
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 91a711977033168e7a7abe06637558169ee7ab7d
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Wed Jul 7 18:07:38 2021 -0700

    hexagon: handle {,SOFT}IRQENTRY_TEXT in linker script

    [ Upstream commit 6fef087d0d37ba7dba8f3d75566eb4c256cd6742 ]

    Patch series "hexagon: Fix build error with CONFIG_STACKDEPOT and select
CONFIG_ARCH_WANT_LD_ORPHAN_WARN".

    This series fixes an error with ARCH=hexagon that was pointed out by the
    patch "mm/slub: use stackdepot to save stack trace in objects".

    The first patch fixes that error by handling the '.irqentry.text' and
    '.softirqentry.text' sections.

    The second patch switches Hexagon over to the common DISCARDS macro, which
    should have been done when Hexagon was merged into the tree to match
    commit 023bf6f1b8bf ("linker script: unify usage of discard definition").

    The third patch selects CONFIG_ARCH_WANT_LD_ORPHAN_WARN so that something
    like this does not happen again.

    This patch (of 3):

    Patch "mm/slub: use stackdepot to save stack trace in objects" in -mm
    selects CONFIG_STACKDEPOT when CONFIG_STACKTRACE_SUPPORT is selected and
    CONFIG_STACKDEPOT requires IRQENTRY_TEXT and SOFTIRQENTRY_TEXT to be
    handled after commit 505a0ef15f96 ("kasan: stackdepot: move
    filter_irq_stacks() to stackdepot.c") due to the use of the
    __{,soft}irqentry_text_{start,end} section symbols.  If those sections are
    not handled, the build is broken.

    $ make ARCH=hexagon CROSS_COMPILE=hexagon-linux- LLVM=1 LLVM_IAS=1 defconfig all
    ...
    ld.lld: error: undefined symbol: __irqentry_text_start
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a

    ld.lld: error: undefined symbol: __irqentry_text_end
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a

    ld.lld: error: undefined symbol: __softirqentry_text_start
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a

    ld.lld: error: undefined symbol: __softirqentry_text_end
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a
    >>> referenced by stackdepot.c
    >>>               stackdepot.o:(filter_irq_stacks) in archive lib/built-in.a
    ...
```

```
    Add these sections to the Hexagon linker script so the build continues to
    work.  ld.lld's orphan section warning would have caught this prior to the
    -mm commit mentioned above:

    ld.lld: warning: kernel/built-in.a(softirq.o):(.softirqentry.text) is being placed in
'.softirqentry.text'
    ld.lld: warning: kernel/built-in.a(softirq.o):(.softirqentry.text) is being placed in
'.softirqentry.text'
    ld.lld: warning: kernel/built-in.a(softirq.o):(.softirqentry.text) is being placed in
'.softirqentry.text'

    Link: https://lkml.kernel.org/r/20210521011239.1332345-1-nathan@kernel.org
    Link: https://lkml.kernel.org/r/20210521011239.1332345-2-nathan@kernel.org
    Link: https://github.com/ClangBuiltLinux/linux/issues/1381
    Fixes: 505a0ef15f96 ("kasan: stackdepot: move filter_irq_stacks() to stackdepot.c")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Acked-by: Brian Cain <bcain@codeaurora.org>
    Cc: Oliver Glitta <glittao@gmail.com>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Cc: David Rientjes <rientjes@google.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5d3913911f12b1cc35471c7a357c7cb0668f8e5
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Sat Jul 3 14:34:20 2021 -0400

    NFSv4/pNFS: Don't call _nfs4_pnfs_v3_ds_connect multiple times

    [ Upstream commit f46f84931a0aa344678efe412d4b071d84d8a805 ]

    After we grab the lock in nfs4_pnfs_ds_connect(), there is no check for
    whether or not ds->ds_clp has already been initialised, so we can end up
    adding the same transports multiple times.

    Fixes: fc821d59209d ("pnfs/NFSv4.1: Add multipath capabilities to pNFS flexfiles servers
over NFSv3")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a43389b7df966d23d595ba34ed43f8c322562f92
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Fri Jul 2 19:48:41 2021 -0400

    NFSv4/pnfs: Fix layoutget behaviour after invalidation

    [ Upstream commit 0b77f97a7e42adc72bd566ff8cb733ea426f74f6 ]

    If the layout gets invalidated, we should wait for any outstanding
    layoutget requests for that layout to complete, and we should resend
    them only after re-establishing the layout stateid.

    Fixes: d29b468da4f9 ("pNFS/NFSv4: Improve rejection of out-of-order layouts")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f90295b4b64189ac12c90615fcf8f1356ac2fa8b
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Fri Jul 2 16:37:15 2021 -0400

    NFSv4/pnfs: Fix the layout barrier update
```

```
    [ Upstream commit aa95edf309ef31e2df4a37ebf0e5c2ca2a6772ab ]

    If we have multiple outstanding layoutget requests, the current code to
    update the layout barrier assumes that the outstanding layout stateids
    are updated in order. That's not necessarily the case.

    Instead of using the value of lo->plh_outstanding as a guesstimate for
    the window of values we need to accept, just wait to update the window
    until we're processing the last one. The intention here is just to
    ensure that we don't process 2^31 seqid updates without also updating
    the barrier.

    Fixes: 1bcf34fdac5f ("pNFS/NFSv4: Update the layout barrier when we schedule a
layoutreturn")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65d105e0014c171381aef276192c0a61604be961
Author: Dave Wysochanski <dwysocha@redhat.com>
Date:   Tue Jun 29 13:13:57 2021 -0400

    NFS: Fix fscache read from NFS after cache error

    [ Upstream commit ba512c1bc3232124567a59a3995c773dc79716e8 ]

    Earlier commits refactored some NFS read code and removed
    nfs_readpage_async(), but neglected to properly fixup
    nfs_readpage_from_fscache_complete().  The code path is
    only hit when something unusual occurs with the cachefiles
    backing filesystem, such as an IO error or while a cookie
    is being invalidated.

    Mark page with PG_checked if fscache IO completes in error,
    unlock the page, and let the VM decide to re-issue based on
    PG_uptodate.  When the VM reissues the readpage, PG_checked
    allows us to skip over fscache and read from the server.

    Link: https://marc.info/?l=linux-nfs&m=162498209518739
    Fixes: 1e83b173b266 ("NFS: Add nfs_pageio_complete_read() and remove
nfs_readpage_async()")
    Signed-off-by: Dave Wysochanski <dwysocha@redhat.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9d94a3f255da41c6a3d17a2378b7e5f3a64f6186
Author: Dave Wysochanski <dwysocha@redhat.com>
Date:   Tue Jun 29 05:11:28 2021 -0400

    NFS: Ensure nfs_readpage returns promptly when internal error occurs

    commit e0340f16a08d031de54ed91d26f57c9a966a776a upstream.

    A previous refactoring of nfs_readpage() might end up calling
    wait_on_page_locked_killable() even if readpage_async_filler() failed
    with an internal error and pg_error was non-zero (for example, if
    nfs_create_request() failed).  In the case of an internal error,
    skip over wait_on_page_locked_killable() as this is only needed
    when the read is sent and an error occurs during completion handling.

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

commit 382f21675997ce101f13d7d635b3fbc5aceeed11
```

```
Author: David Hildenbrand <david@redhat.com>
Date:   Wed Jun 2 20:57:14 2021 +0200

    virtio-mem: don't read big block size in Sub Block Mode

    [ Upstream commit 500817bf5e110ad9b7138bc582971bb7ee77d6f7 ]

    We are reading a Big Block Mode value while in Sub Block Mode
    when initializing. Fortunately, vm->bbm.bb_size maps to some counter
    in the vm->sbm.mb_count array, which is 0 at that point in time.

    No harm done; still, this was unintended and is not future-proof.

    Fixes: 4ba50cd3355d ("virtio-mem: Big Block Mode (BBM) memory hotplug")
    Signed-off-by: David Hildenbrand <david@redhat.com>
    Link: https://lore.kernel.org/r/20210602185720.31821-2-david@redhat.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e0bfe5edb790d2472801794224213739 9e01485
Author: Eli Cohen <elic@nvidia.com>
Date:   Sun Jun 6 08:31:28 2021 +0300

    vdpa/mlx5: Clear vq ready indication upon device reset

    [ Upstream commit e3aadf2e1614174dc81d52cbb9dabb77913b11c6 ]

    After device reset, the virtqueues are not ready so clear the ready
    field.

    Failing to do so can result in virtio_vdpa failing to load if the device
    was previously used by vhost_vdpa and the old values are ready.
    virtio_vdpa expects to find VQs in "not ready" state.

    Fixes: 1a86b377aa21 ("vdpa/mlx5: Add VDPA driver for supported mlx5 devices")
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20210606053128.170399-1-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eaa0eb7020f7c71ae37b1bc51ffc435e0cfcdb69
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Wed Jul 7 15:40:51 2021 +0800

    ALSA: isa: Fix error return code in snd_cmi8330_probe()

    [ Upstream commit 31028cbed26a8afa25533a10425ffa2ab794c76c ]

    When 'SB_HW_16' check fails, the error code -ENODEV instead of 0 should be
    returned, which is the same as that returned when 'WSS_HW_CMI8330' check
    fails.

    Fixes: 43bcd973d6d0 ("[ALSA] Add snd_card_set_generic_dev() call to ISA drivers")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Link: https://lore.kernel.org/r/20210707074051.2663-1-thunder.leizhen@huawei.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 650e6f383a6eb40f7c0a010982a74ab4b6893870
Author: J. Bruce Fields <bfields@redhat.com>
Date:   Thu Jul 1 20:06:56 2021 -0400

    nfsd: fix NULL dereference in nfs3svc_encode_getaclres
```

```
    [ Upstream commit ab1016d39cc052064e32f25ad18ef8767a0ee3b8 ]

    In error cases the dentry may be NULL.

    Before 20798dfe249a, the encoder also checked dentry and
    d_really_is_positive(dentry), but that looks like overkill to me--zero
    status should be enough to guarantee a positive dentry.

    This isn't the first time we've seen an error-case NULL dereference
    hidden in the initialization of a local variable in an xdr encoder.  But
    I went back through the other recent rewrites and didn't spot any
    similar bugs.

    Reported-by: JianHong Yin <jiyin@redhat.com>
    Reviewed-by: Chuck Lever III <chuck.lever@oracle.com>
    Fixes: 20798dfe249a ("NFSD: Update the NFSv3 GETACL result encoder...")
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7605bff387a9972038b217b6c60998778dbae931
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Fri Jun 25 11:12:49 2021 -0400

    NFSD: Prevent a possible oops in the nfs_dirent() tracepoint

    [ Upstream commit 7b08cf62b1239a4322427d677ea9363f0ab677c6 ]

    The double copy of the string is a mistake, plus __assign_str()
    uses strlen(), which is wrong to do on a string that isn't
    guaranteed to be NUL-terminated.

    Fixes: 6019ce0742ca ("NFSD: Add a tracepoint to record directory entry encoding")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8a8bb8dca84a8f590f0293b5d42c62e370ab931
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Thu Jun 17 19:26:52 2021 -0400

    nfsd: Reduce contention for the nfsd_file nf_rwsem

    [ Upstream commit 474bc334698df98ce07c890f1898c7e7f389b0c7 ]

    When flushing out the unstable file writes as part of a COMMIT call, try
    to perform most of of the data writes and waits outside the semaphore.

    This means that if the client is sending the COMMIT as part of a memory
    reclaim operation, then it can continue performing I/O, with contention
    for the lock occurring only once the data sync is finished.

    Fixes: 5011af4c698a ("nfsd: Fix stable writes")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
     Tested-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1e8c6f23b183f40fcb4e8fbc4e558b6f125db967
Author: J. Bruce Fields <bfields@redhat.com>
Date:   Tue May 25 14:53:44 2021 -0400

    nfsd: move fsnotify on client creation outside spinlock

    [ Upstream commit 934bd07fae7e55232845f909f78873ab8678ca74 ]
```

```
    This was causing a "sleeping function called from invalid context"
    warning.

    I don't think we need the set_and_test_bit() here; clients move from
    unconfirmed to confirmed only once, under the client_lock.

    The (conf == unconf) is a way to check whether we're in that confirming
    case, hopefully that's not too obscure.

    Fixes: 472d155a0631 "nfsd: report client confirmation status in "info" file"
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ca9c3d31c52d585ed46ba8c2fbd77d4d02d389c2
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Fri May 14 15:55:48 2021 -0400

    NFSD: Add nfsd_clid_confirmed tracepoint

    [ Upstream commit 7e3b32ace6094aadfa2e1e54ca4c6bbfd07646af ]

    This replaces a dprintk call site in order to get greater visibility
    on when client IDs are confirmed or re-used. Simple example:

                nfsd-995   [000]   126.622975: nfsd_compound:        xid=0x3a34e2b1 opcnt=1
                nfsd-995   [000]   126.623005: nfsd_cb_args:         addr=192.168.2.51:45901
client 60958e3b:9213ef0e prog=1073741824 ident=1
                nfsd-995   [000]   126.623007: nfsd_compound_status: op=1/1 OP_SETCLIENTID
status=0
                nfsd-996   [001]   126.623142: nfsd_compound:        xid=0x3b34e2b1 opcnt=1
      >>>>      nfsd-996   [001]   126.623146: nfsd_clid_confirmed:  client
60958e3b:9213ef0e
                nfsd-996   [001]   126.623148: nfsd_cb_probe:        addr=192.168.2.51:45901
client 60958e3b:9213ef0e state=UNKNOWN
                nfsd-996   [001]   126.623154: nfsd_compound_status: op=1/1
OP_SETCLIENTID_CONFIRM status=0

    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d435b6d94b05dcfd836d758a63145aa566618e2
Author: Naveen N. Rao <naveen.n.rao@linux.vnet.ibm.com>
Date:   Thu Jul 1 20:38:58 2021 +0530

    powerpc/bpf: Fix detecting BPF atomic instructions

    [ Upstream commit 419ac821766cbdb9fd85872bb3f1a589df05c94c ]

    Commit 91c960b0056672 ("bpf: Rename BPF_XADD and prepare to encode other
    atomics in .imm") converted BPF_XADD to BPF_ATOMIC and added a way to
    distinguish instructions based on the immediate field. Existing JIT
    implementations were updated to check for the immediate field and to
    reject programs utilizing anything more than BPF_ADD (such as BPF_FETCH)
    in the immediate field.

    However, the check added to powerpc64 JIT did not look at the correct
    BPF instruction. Due to this, such programs would be accepted and
    incorrectly JIT'ed resulting in soft lockups, as seen with the atomic
    bounds test. Fix this by looking at the correct immediate value.

    Fixes: 91c960b0056672 ("bpf: Rename BPF_XADD and prepare to encode other atomics in
.imm")
    Reported-by: Jiri Olsa <jolsa@redhat.com>
```

```
    Signed-off-by: Naveen N. Rao <naveen.n.rao@linux.vnet.ibm.com>
    Tested-by: Jiri Olsa <jolsa@redhat.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/4117b430ffaa8cd7af042496f87fd7539e4f17fd.1625145429.git.naveen.n.r
ao@linux.vnet.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc6b8e2cd87e894613e331a636ac93168712466e
Author: Maurizio Lombardi <mlombard@redhat.com>
Date:   Fri Jul 2 10:11:21 2021 +0200

    nvme-tcp: can't set sk_user_data without write_lock

    [ Upstream commit 0755d3be2d9bb6ea38598ccd30d6bbaa1a5c3a50 ]

    The sk_user_data pointer is supposed to be modified only while
    holding the write_lock "sk_callback_lock", otherwise
    we could race with other threads and crash the kernel.

    we can't take the write_lock in nvmet_tcp_state_change()
    because it would cause a deadlock, but the release_work queue
    will set the pointer to NULL later so we can simply remove
    the assignment.

    Fixes: b5332a9f3f3d ("nvmet-tcp: fix incorrect locking in state_change sk callback")

    Signed-off-by: Maurizio Lombardi <mlombard@redhat.com>
    Reviewed-by: Sagi Grimberg <sagi@grimberg.me>
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c01f6cae4b1b813f18777170e7c25940d47c967b
Author: Michael S. Tsirkin <mst@redhat.com>
Date:   Tue Apr 13 01:35:26 2021 -0400

    virtio_net: move tx vq operation under tx queue lock

    [ Upstream commit 5a2f966d0f3fa0ef6dada7ab9eda74cacee96b8a ]

    It's unsafe to operate a vq from multiple threads.
    Unfortunately this is exactly what we do when invoking
    clean tx poll from rx napi.
    Same happens with napi-tx even without the
    opportunistic cleaning from the receive interrupt: that races
    with processing the vq in start_xmit.

    As a fix move everything that deals with the vq to under tx lock.

    Fixes: b92f1e6751a6 ("virtio-net: transmit napi")
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70899667e59a70c95f7dbc78403a626ab4cf0710
Author: Eli Cohen <elic@nvidia.com>
Date:   Sun Jun 6 08:31:50 2021 +0300

    vdp/mlx5: Fix setting the correct dma_device

    [ Upstream commit 7d23dcdf213c2e5f097eb7eec3148c26eb01d59f ]

    Before SF support was introduced, the DMA device was equal to
    mdev->device which was in essence equal to pdev->dev.

    With SF introduction this is no longer true. It has already been
```

```
    handled for vhost_vdpa since the reference to the dma device can from
    within mlx5_vdpa. With virtio_vdpa this broke. To fix this we set the
    real dma device when initializing the device.

    In addition, for the sake of consistency, previous references in the
    code to the dma device are changed to vdev->dma_dev.

    Fixes: d13a15d544ce5 ("vdpa/mlx5: Use the correct dma device when registering memory")
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20210606053150.170489-1-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e654dbc9a6e2efc2845a11093040563b7ee57150
Author: Eli Cohen <elic@nvidia.com>
Date:   Sun May 30 12:03:49 2021 +0300

    vdpa/mlx5: Fix possible failure in umem size calculation

    [ Upstream commit 71ab6a7cfbae27f86a3901daab10bfe13b3a1e3a ]

    umem size is a 32 bit unsigned value so assigning it to an int could
    cause false failures. Set the calculated value inside the function and
    modify function name to reflect the fact it updates the size.

    This bug was found during code review but never had real impact to this
    date.

    Fixes: 1a86b377aa21 ("vdpa/mlx5: Add VDPA driver for supported mlx5 devices")
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20210530090349.8360-1-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 965b080f4e6d2ec1f0a3dc2aa1a06ee94c90d603
Author: Eli Cohen <elic@nvidia.com>
Date:   Sun May 30 12:03:17 2021 +0300

    vdpa/mlx5: Fix umem sizes assignments on VQ create

    [ Upstream commit e3011776af16caf423f2c36d0047acd624c274fa ]

    Fix copy paste bug assigning umem1 size to umem2 and umem3. The issue
    was discovered when trying to use a 1:1 MR that covers the entire
    address space where firmware complained that provided sizes are not
    large enough. 1:1 MRs are required to support virtio_vdpa.

    Fixes: 1a86b377aa21 ("vdpa/mlx5: Add VDPA driver for supported mlx5 devices")
    Signed-off-by: Eli Cohen <elic@nvidia.com>
    Link: https://lore.kernel.org/r/20210530090317.8284-1-elic@nvidia.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9dca259fbb420a071597bc0ebfd8e30a53190d75
Author: Jason Wang <jasowang@redhat.com>
Date:   Thu Jun 24 11:59:39 2021 +0800

    vp_vdpa: correct the return value when fail to map notification

    [ Upstream commit 94e48d6aafef23143f92eadd010c505c49487576 ]

    We forget to assign a error value when we fail to map the notification
```

```
    during prove. This patch fixes it.

    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@oracle.com>
    Fixes: 11d8ffed00b23 ("vp_vdpa: switch to use vp_modern_map_vq_notify()")
    Signed-off-by: Jason Wang <jasowang@redhat.com>
    Link: https://lore.kernel.org/r/20210624035939.26618-1-jasowang@redhat.com
    Reviewed-by: Stefano Garzarella <sgarzare@redhat.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c31fab84b23c779eee2ff2d45ea7316fb31b223b
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Sat Jun 26 21:40:46 2021 +0100

    io_uring: remove not needed PF_EXITING check

    [ Upstream commit e5dc480d4ed9884274e95c757fa2d2e9cc1047ee ]

    Since cancellation got moved before exit_signals(), there is no one left
    who can call io_run_task_work() with PF_EXIING set, so remove the check.
    Note that __io_req_task_submit() still needs a similar check.

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/f7f305ececb1e6044ea649fb983ca754805bb884.1624739600.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dafcb53b550b0ae8f442d3e910b928424803f123
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Thu Jun 17 18:14:06 2021 +0100

    io_uring: inline __tctx_task_work()

    [ Upstream commit 3f18407dc6f2db0968daaa36c39a772c2c9f8ea7 ]

    Inline __tctx_task_work() into tctx_task_work() in preparation for
    further optimisations.

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/f9c05c4bc9763af7bd8e25ebc3c5f7b6f69148f8.1623949695.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7511a4f524f3dde008ccc72a592dcbde5b59de38
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Thu Jun 17 18:14:01 2021 +0100

    io_uring: move creds from io-wq work to io_kiocb

    [ Upstream commit c10d1f986b4e2a906862148c77a97f186cc08b9e ]

    io-wq now doesn't have anything to do with creds now, so move ->creds
    from struct io_wq_work into request (aka struct io_kiocb).

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/8520c72ab8b8f4b96db12a228a2ab4c094ae64e1.1623949695.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 16eadf8979564c53a5a9aaa0203ee87a034cc59d
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Sun May 16 22:58:12 2021 +0100

    io_uring: don't bounce submit_state cachelines

    [ Upstream commit d0acdee296d42e700c16271d9f95085a9c897a53 ]

    struct io_submit_state contains struct io_comp_state and so
    locked_free_*, that renders cachelines around ->locked_free* being
    invalidated on most non-inline completions, that may terrorise caches if
    submissions and completions are done by different tasks.

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/290cb5412b76892e8631978ee8ab9db0c6290dd5.1621201931.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bea4688b9e5928be269135549b0b65c3ae2bc8be
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Sun May 16 22:58:06 2021 +0100

    io_uring: shuffle rarely used ctx fields

    [ Upstream commit b986af7e2df4f0871367c397ba61a542f37c0ab3 ]

    There is a bunch of scattered around ctx fields that are almost never
    used, e.g. only on ring exit, plunge them to the end, better locality,
    better aesthetically.

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/782ff94b00355923eae757d58b1a47821b5b46d4.1621201931.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f564954855fcc247c953be50fd9d480110e73cad
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Sun May 16 22:58:04 2021 +0100

    io_uring: get rid of files in exit cancel

    [ Upstream commit 3dd0c97a9e011b11ce6bd245bacf58c57f6f7875 ]

    We don't match against files on cancellation anymore, so no need to drag
    around files_struct anymore, just pass a flag telling whether only
    inflight or all requests should be killed.

    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/7bfc5409a78f8e2d6b27dec3293ec2d248677348.1621201931.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04c529ed0e7f40848869e0dc8c4f9b7c51d9133e
Author: Christoph Hellwig <hch@lst.de>
Date:   Thu Jul 1 10:16:37 2021 +0200

    block: grab a device refcount in disk_uevent
```

```
    [ Upstream commit 498dcc13fd6463de29b94e160f40ed04d5477cd8 ]

    Sending uevents requires the struct device to be alive.  To
    ensure that grab the device refcount instead of just an inode
    reference.

    Fixes: bc359d03c7ec ("block: add a disk_uevent helper")
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20210701081638.246552-2-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ccdf7e073170886bc370c613e269de610a794c4a
Author: Viresh Kumar <viresh.kumar@linaro.org>
Date:   Tue Jun 15 14:27:50 2021 +0530

    arch_topology: Avoid use-after-free for scale_freq_data

    [ Upstream commit 83150f5d05f065fb5c12c612f119015cabdcc124 ]

    Currently topology_scale_freq_tick() (which gets called from
    scheduler_tick()) may end up using a pointer to "struct
    scale_freq_data", which was previously cleared by
    topology_clear_scale_freq_source(), as there is no protection in place
    here. The users of topology_clear_scale_freq_source() though needs a
    guarantee that the previously cleared scale_freq_data isn't used
    anymore, so they can free the related resources.

    Since topology_scale_freq_tick() is called from scheduler tick, we don't
    want to add locking in there. Use the RCU update mechanism instead
    (which is already used by the scheduler's utilization update path) to
    guarantee race free updates here.

    synchronize_rcu() makes sure that all RCU critical sections that started
    before it is called, will finish before it returns. And so the callers
    of topology_clear_scale_freq_source() don't need to worry about their
    callback getting called anymore.

    Cc: Paul E. McKenney <paulmck@kernel.org>
    Fixes: 01e055c120a4 ("arch_topology: Allow multiple entities to provide
sched_freq_tick() callback")
    Tested-by: Vincent Guittot <vincent.guittot@linaro.org>
    Reviewed-by: Ionela Voinescu <ionela.voinescu@arm.com>
    Tested-by: Qian Cai <quic_qiancai@quicinc.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d6c153a6ae03f49dc1ee6f2593e9d2dcd33c8c0e
Author: Jon Hunter <jonathanh@nvidia.com>
Date:   Fri Jun 18 17:02:19 2021 +0100

    PCI: tegra194: Fix tegra_pcie_ep_raise_msi_irq() ill-defined shift

    [ Upstream commit f67092eff2bd40650aad54a1a1910160f41d864a ]

    tegra_pcie_ep_raise_msi_irq() shifted a signed 32-bit value left by 31
    bits.  The behavior of this is implementation-defined.

    Replace the shift by BIT(), which is well-defined.

    Found by cppcheck:

      $ cppcheck --enable=all drivers/pci/controller/dwc/pcie-tegra194.c
      Checking drivers/pci/controller/dwc/pcie-tegra194.c ...
```

```
      drivers/pci/controller/dwc/pcie-tegra194.c:1829:23: portability: Shifting signed 32-
 bit value by 31 bits is implementation-defined behaviour. See condition at line 1826.
 [shiftTooManyBitsSigned]

      appl_writel(pcie, (1 << irq), APPL_MSI_CTRL_1);
                        ^


    [bhelgaas: commit log]
    Link: https://lore.kernel.org/r/20210618160219.303092-1-jonathanh@nvidia.com
    Fixes: c57247f940e8 ("PCI: tegra: Add support for PCIe endpoint mode in Tegra194")
    Signed-off-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63798e31f5888e322491863f88e9c03c435ce116
Author: Jaegeuk Kim <jaegeuk@kernel.org>
Date:   Tue Jun 22 12:56:44 2021 -0700

    f2fs: remove false alarm on iget failure during GC

    [ Upstream commit 132e3209789c647e37dc398ef36af4de13f104b4 ]

    This patch removes setting SBI_NEED_FSCK when GC gets an error on f2fs_iget,
    since f2fs_iget can give ENOMEM and others by race condition.
    If we set this critical fsck flag, we'll get EIO during fsync via the below
    code path.

    In f2fs_inplace_write_data(),

            if (is_sbi_flag_set(sbi, SBI_NEED_FSCK) || f2fs_cp_error(sbi)) {
                    err = -EIO;
                    goto drop_bio;
            }

    Fixes: 9557727876674 ("f2fs: drop inplace IO if fs status is abnormal")
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e882298f7c3ef96d41ddd8aec0df0a65b59577a5
Author: Scott Mayhew <smayhew@redhat.com>
Date:   Tue Jun 22 08:11:59 2021 -0400

    nfs: update has_sec_mnt_opts after cloning lsm options from parent

    [ Upstream commit eae00c5d6e48ccb2d78ae5873743d7d1a572951b ]

    After calling security_sb_clone_mnt_opts() in nfs_get_root(), it's
    necessary to copy the value of has_sec_mnt_opts from the cloned
    super_block's nfs_server.  Otherwise, calls to nfs_compare_super()
    using this super_block may not return the correct result, leading to
    mount failures.

    For example, mounting an nfs server with the following in /etc/exports:
    /export *(rw,insecure,crossmnt,no_root_squash,security_label)
    and having /export/scratch on a separate block device.

    mount -o v4.2,context=system_u:object_r:root_t:s0 server:/export/test /mnt/test
    mount -o v4.2,context=system_u:object_r:swapfile_t:s0 server:/export/scratch
 /mnt/scratch

    The second mount would fail with "mount.nfs: /mnt/scratch is busy or
    already mounted or sharecache fail" and "SELinux: mount invalid.  Same
    superblock, different security settings for..." would appear in the
```

```
    syslog.

    Also while we're in there, replace several instances of "NFS_SB(s)"
    with "server", which was already declared at the top of the
    nfs_get_root().

    Fixes: ec1ade6a0448 ("nfs: account for selinux security context when deciding to share
superblock")
    Signed-off-by: Scott Mayhew <smayhew@redhat.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 582223e621e1c08446d1f95d9356d40d3664a4ea
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Tue May 25 08:35:27 2021 +0200

    pwm: imx1: Don't disable clocks at device remove time

    [ Upstream commit 1bc6ea31cb41d50302a3c9b401964cf0a88d41f9 ]

    The .remove() callback disables clocks that were not enabled in
    .probe(). So just probing and then unbinding the driver results in a clk
    enable imbalance.

    So just drop the call to disable the clocks. (Which BTW was also in the
    wrong order because the call makes the PWM unfunctional and so should
    have come only after pwmchip_remove()).

    Fixes: 9f4c8f9607c3 ("pwm: imx: Add ipg clock operation")
    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 18cfc69d561a4a61ad6163a6ce87b48c8d1cab6b
Author: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
Date:   Wed Jan 6 14:55:40 2021 +0100

    PCI: intel-gw: Fix INTx enable

    [ Upstream commit 655832d12f2251e04031294f547c86935a0a126d ]

    The legacy PCI interrupt lines need to be enabled using PCIE_APP_IRNEN bits
    13 (INTA), 14 (INTB), 15 (INTC) and 16 (INTD). The old code however was
    taking (for example) "13" as raw value instead of taking BIT(13).  Define
    the legacy PCI interrupt bits using the BIT() macro and then use these in
    PCIE_APP_IRN_INT.

    Link: https://lore.kernel.org/r/20210106135540.48420-1-
martin.blumenstingl@googlemail.com
    Fixes: ed22aaaede44 ("PCI: dwc: intel: PCIe RC controller driver")
    Signed-off-by: Martin Blumenstingl <martin.blumenstingl@googlemail.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Rahul Tanwar <rtanwar@maxlinear.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29e9a35b400b978aa541f9f97698af1651bed05c
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Jun 23 14:01:35 2021 +0200

    x86/fpu: Limit xstate copy size in xstateregs_set()

    [ Upstream commit 07d6688b22e09be465652cf2da0da6bf86154df6 ]

    If the count argument is larger than the xstate size, this will happily
```

```
    copy beyond the end of xstate.

    Fixes: 91c3dba7dbc1 ("x86/fpu/xstate: Fix PTRACE frames for XSAVES")
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Andy Lutomirski <luto@kernel.org>
    Reviewed-by: Borislav Petkov <bp@suse.de>
    Link: https://lkml.kernel.org/r/20210623121452.120741557@linutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a04e670b07df36612522731b1de2b013beade21c
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Jun 23 14:01:28 2021 +0200

    x86/fpu: Fix copy_xstate_to_kernel() gap handling

    [ Upstream commit 9625895011d130033d1bc7aac0d77a9bf68ff8a6 ]

    The gap handling in copy_xstate_to_kernel() is wrong when XSAVES is in
    use.

    Using init_fpstate for copying the init state of features which are
    not set in the xstate header is only correct for the legacy area, but
    not for the extended features area because when XSAVES is in use then
    init_fpstate is in compacted form which means the xstate offsets which
    are used to copy from init_fpstate are not valid.

    Fortunately, this is not a real problem today because all extended
    features in use have an all-zeros init state, but it is wrong
    nevertheless and with a potentially dynamically sized init_fpstate this
    would result in an access outside of the init_fpstate.

    Fix this by keeping track of the last copied state in the target buffer and
    explicitly zero it when there is a feature or alignment gap.

    Use the compacted offset when accessing the extended feature space in
    init_fpstate.

    As this is not a functional issue on older kernels this is intentionally
    not tagged for stable.

    Fixes: b8be15d58806 ("x86/fpu/xstate: Re-enable XSAVES")
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Borislav Petkov <bp@suse.de>
    Link: https://lkml.kernel.org/r/20210623121451.294282032@linutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c615921c48c98ea4a29cf297fce850abc6f436b3
Author: Chao Yu <chao@kernel.org>
Date:   Tue Jun 8 07:31:22 2021 +0800

    f2fs: fix to avoid adding tab before doc section

    [ Upstream commit 3c16dc40aab84bab9cf54c2b61a458bb86b180c3 ]

    Otherwise whole section after tab will be invisible in compiled
    html format document.

    Cc: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Fixes: 89272ca1102e ("docs: filesystems: convert f2fs.txt to ReST")
    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 7bc134df92a2a8187a06d74293e91f6dfd86fbc4
Author: Sandor Bodo-Merle <sbodomerle@gmail.com>
Date:   Tue Jun 22 17:26:30 2021 +0200

    PCI: iproc: Support multi-MSI only on uniprocessor kernel

    [ Upstream commit 2dc0a201d0f59e6818ef443609f0850a32910844 ]

    The interrupt affinity scheme used by this driver is incompatible with
    multi-MSI as it implies moving the doorbell address to that of another MSI
    group.  This isn't possible for multi-MSI, as all the MSIs must have the
    same doorbell address. As such it is restricted to systems with a single
    CPU.

    Link: https://lore.kernel.org/r/20210622152630.40842-2-sbodomerle@gmail.com
    Fixes: fc54bae28818 ("PCI: iproc: Allow allocation of multiple MSIs")
    Reported-by: Marc Zyngier <maz@kernel.org>
    Signed-off-by: Sandor Bodo-Merle <sbodomerle@gmail.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Acked-by: Marc Zyngier <maz@kernel.org>
    Acked-by: Pali Rohár <pali@kernel.org>
    Acked-by: Ray Jui <ray.jui@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5ac5596c3b7c6218623e09a70e5b82ff8584ede4
Author: Sandor Bodo-Merle <sbodomerle@gmail.com>
Date:   Tue Jun 22 17:26:29 2021 +0200

    PCI: iproc: Fix multi-MSI base vector number allocation

    [ Upstream commit e673d697b9a234fc3544ac240e173cef8c82b349 ]

    Commit fc54bae28818 ("PCI: iproc: Allow allocation of multiple MSIs")
    introduced multi-MSI support with a broken allocation mechanism (it failed
    to reserve the proper number of bits from the inner domain).  Natural
    alignment of the base vector number was also not guaranteed.

    Link: https://lore.kernel.org/r/20210622152630.40842-1-sbodomerle@gmail.com
    Fixes: fc54bae28818 ("PCI: iproc: Allow allocation of multiple MSIs")
    Reported-by: Pali Rohár <pali@kernel.org>
    Signed-off-by: Sandor Bodo-Merle <sbodomerle@gmail.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Acked-by: Marc Zyngier <maz@kernel.org>
    Acked-by: Pali Rohár <pali@kernel.org>
    Acked-by: Ray Jui <ray.jui@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8f02865de3a3358a3ccaddb4eb3017fa24e05c54
Author: Zhihao Cheng <chengzhihao1@huawei.com>
Date:   Fri Jun 18 16:11:03 2021 +0800

    ubifs: Set/Clear I_LINKABLE under i_lock for whiteout inode

    [ Upstream commit a801fcfeef96702fa3f9b22ad56c5eb1989d9221 ]

    xfstests-generic/476 reports a warning message as below:

    WARNING: CPU: 2 PID: 30347 at fs/inode.c:361 inc_nlink+0x52/0x70
    Call Trace:
      do_rename+0x502/0xd40 [ubifs]
      ubifs_rename+0x8b/0x180 [ubifs]
      vfs_rename+0x476/0x1080
      do_renameat2+0x67c/0x7b0
      __x64_sys_renameat2+0x6e/0x90
      do_syscall_64+0x66/0xe0
```

```
       entry_SYSCALL_64_after_hwframe+0x44/0xae

    Following race case can cause this:
             rename_whiteout(Thread 1)             wb_workfn(Thread 2)
    ubifs_rename
      do_rename
                                              __writeback_single_inode
                                                spin_lock(&inode->i_lock)
        whiteout->i_state |= I_LINKABLE
                                                inode->i_state &= ~dirty;
    ---- How race happens on i_state:
        (tmp = whiteout->i_state | I_LINKABLE)
                                               (tmp = inode->i_state & ~dirty)
        (whiteout->i_state = tmp)
                                               (inode->i_state = tmp)
    ----
                                                spin_unlock(&inode->i_lock)
        inc_nlink(whiteout)
        WARN_ON(!(inode->i_state & I_LINKABLE)) !!!

    Fix to add i_lock to avoid i_state update race condition.

    Fixes: 9e0a1fff8db56ea ("ubifs: Implement RENAME_WHITEOUT")
    Signed-off-by: Zhihao Cheng <chengzhihao1@huawei.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 687cf32865b2d6960214bce523f2afac58dd3cd2
Author: Gao Xiang <hsiangkao@linux.alibaba.com>
Date:   Fri Jun 18 12:20:55 2021 +0800

    nfs: fix acl memory leak of posix_acl_create()

    [ Upstream commit 1fcb6fcd74a222d9ead54d405842fc763bb86262 ]

    When looking into another nfs xfstests report, I found acl and
    default_acl in nfs3_proc_create() and nfs3_proc_mknod() error
    paths are possibly leaked. Fix them in advance.

    Fixes: 013cdf1088d7 ("nfs: use generic posix ACL infrastructure for v3 Posix ACLs")
    Cc: Trond Myklebust <trond.myklebust@hammerspace.com>
    Cc: Anna Schumaker <anna.schumaker@netapp.com>
    Cc: Christoph Hellwig <hch@infradead.org>
    Cc: Joseph Qi <joseph.qi@linux.alibaba.com>
    Signed-off-by: Gao Xiang <hsiangkao@linux.alibaba.com>
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7cedf369663f15de84e28e934de9043ebfc3dfcf
Author: NeilBrown <neilb@suse.de>
Date:   Tue Jun 15 11:18:38 2021 +1000

    SUNRPC: prevent port reuse on transports which don't request it.

    [ Upstream commit bc1c56e9bbe92766d017efb5f0a0c71f80da5570 ]

    If an RPC client is created without RPC_CLNT_CREATE_REUSEPORT, it should
    not reuse the source port when a TCP connection is re-established.
    This is currently implemented by preventing the source port being
    recorded after a successful connection (the call to xs_set_srcport()).

    However the source port is also recorded after a successful bind in xs_bind().
    This may not be needed at all and certainly is not wanted when
    RPC_CLNT_CREATE_REUSEPORT wasn't requested.
```

```
    So avoid that assignment when xprt.reuseport is not set.

    With this change, NFSv4.1 and later mounts use a different port number on
    each connection.  This is helpful with some firewalls which don't cope
    well with port reuse.

    Signed-off-by: NeilBrown <neilb@suse.de>
    Fixes: e6237b6feb37 ("NFSv4.1: Don't rebind to the same source port when reconnecting to
the server")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a96036c0b520d5b90ff63c4251895f13a05e9f4b
Author: Wei Yongjun <weiyongjun1@huawei.com>
Date:   Thu Mar 4 04:59:09 2021 +0000

    watchdog: jz4740: Fix return value check in jz4740_wdt_probe()

    [ Upstream commit 29e85f53fb58b45b9e9276dcdf1f1cb762dd1c9f ]

    In case of error, the function device_node_to_regmap() returns
    ERR_PTR() and never returns NULL. The NULL test in the return
    value check should be replaced with IS_ERR().

    Fixes: 6d532143c915 ("watchdog: jz4740: Use regmap provided by TCU driver")
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Wei Yongjun <weiyongjun1@huawei.com>
    Acked-by: Paul Cercueil <paul@crapouillou.net>
    Link: https://lore.kernel.org/r/20210304045909.945799-1-weiyongjun1@huawei.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07777127a44af1772912eb3217c6cd2a79c6cec7
Author: Tao Ren <rentao.bupt@gmail.com>
Date:   Fri Apr 16 20:42:49 2021 -0700

    watchdog: aspeed: fix hardware timeout calculation

    [ Upstream commit e7dc481c92060f9ce872878b0b7a08c24713a7e5 ]

    Fix hardware timeout calculation in aspeed_wdt_set_timeout function to
    ensure the reload value does not exceed the hardware limit.

    Fixes: efa859f7d786 ("watchdog: Add Aspeed watchdog driver")
    Reported-by: Amithash Prasad <amithash@fb.com>
    Signed-off-by: Tao Ren <rentao.bupt@gmail.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20210417034249.5978-1-rentao.bupt@gmail.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90d61adce9586d2b2e0bf6c021a8c1233e16970f
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:50 2021 +0530

    watchdog: keembay: Removed timeout update in the TO ISR

    [ Upstream commit 3168be5d66ac6c3508a880022f79b5a887865d5d ]

    In the TO ISR removed updating the Timeout value because
    its not serving any purpose as the timer would have already expired
    and the system would be rebooting.
```

```
    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-7-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04693e05d183e50637bb7036b7bbd33bfd2025f0
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:49 2021 +0530

    watchdog: keembay: Remove timeout update in the WDT start function

    [ Upstream commit 9eb25269271c679e8cfcc7df5c0c5e9d0572fc27 ]

    Removed set timeout from the start WDT function. There is a function
    defined to set the timeout. Hence no need to set the timeout again in
    start function as the timeout would have been already updated
    before calling the start/enable.

    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-6-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6479c1a2ae2ce0604378e4f20ba2643a8722dc82
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:48 2021 +0530

    watchdog: keembay: Clear either the TO or TH interrupt bit

    [ Upstream commit 0e36a09faea25f4564d41a0c28938199b605148e ]

    During the interrupt service routine of the TimeOut interrupt and
    the ThresHold interrupt, the respective interrupt clear bit
    have to be cleared and not both.

    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-5-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b8ba05635e6ee3123f037e5cd4e524f474b6762e
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:47 2021 +0530

    watchdog: keembay: Update pretimeout to zero in the TH ISR

    [ Upstream commit 75f6c56dfeec92c53e09a72896547888ac9a27d7 ]

    The pretimeout has to be updated to zero during the ISR of the
    ThresHold interrupt. Else the TH interrupt would be triggerred for
```

```
    every tick until the timeout.

    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-4-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 340590af29d5913feaecb1c3c8c1cf0498801e5a
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:46 2021 +0530

    watchdog: keembay: Upadate WDT pretimeout for every update in timeout

    [ Upstream commit 0f7bfaf10c0abc979220442bae2af4f1f869c41e ]

    The pre-timeout value to be programmed to the register has to be
    calculated and updated for every change in the timeout value.
    Else the threshold time wouldn't be calculated to its
    corresponding timeout.

    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-3-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cfc11b3a8fe58beabb86e8d21cadc2f75d61919
Author: Shruthi Sanil <shruthi.sanil@intel.com>
Date:   Mon May 17 23:19:45 2021 +0530

    watchdog: keembay: Update WDT pre-timeout during the initialization

    [ Upstream commit 29353816300c79cb5157ed2719cc71285c7b77aa ]

    The pretimeout register has a default reset value. Hence
    when a smaller WDT timeout is set which would be lesser than the
    default pretimeout, the system behaves abnormally, starts
    triggering the pretimeout interrupt even when the WDT is
    not enabled, most of the times leading to system crash.
    Hence an update in the pre-timeout is also required for the
    default timeout that is being configured.

    Fixes: fa0f8d51e90d ("watchdog: Add watchdog driver for Intel Keembay Soc")
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Tested-by: Kris Pan <kris.pan@intel.com>
    Signed-off-by: Shruthi Sanil <shruthi.sanil@intel.com>
    Link: https://lore.kernel.org/r/20210517174953.19404-2-shruthi.sanil@intel.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f55dd38490f75e6632eae8f16bfe6b8a0a753e94
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Sat May 8 11:33:13 2021 +0800
```

```
      ubifs: journal: Fix error return code in ubifs_jnl_write_inode()

      [ Upstream commit a2c2a622d41168f9fea2aa3f76b9fbaa88531aac ]

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

      Fixes: 9ca2d7326444 ("ubifs: Limit number of xattrs per inode")
      Reported-by: Hulk Robot <hulkci@huawei.com>
      Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
      Signed-off-by: Richard Weinberger <richard@nod.at>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fe597a0eb351339c18d537792eb467da4379f4ba
  Author: Zhen Lei <thunder.leizhen@huawei.com>
  Date:   Sat May 8 11:22:39 2021 +0800

      um: fix error return code in winch_tramp()

      [ Upstream commit ccf1236ecac476d9d2704866d9a476c86e387971 ]

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

      Fixes: 89df6bfc0405 ("uml: DEBUG_SHIRQ fixes")
      Reported-by: Hulk Robot <hulkci@huawei.com>
      Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
      Acked-By: anton.ivanov@cambridgegreys.com
      Signed-off-by: Richard Weinberger <richard@nod.at>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit a04880c5bb081bed339da63a7c1eebfd4c483689
  Author: Zhen Lei <thunder.leizhen@huawei.com>
  Date:   Sat May 8 11:13:54 2021 +0800

      um: fix error return code in slip_open()

      [ Upstream commit b77e81fbe5f5fb4ad9a61ec80f6d1e30b6da093a ]

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

      Fixes: a3c77c67a443 ("[PATCH] uml: slirp and slip driver cleanups and fixes")
      Reported-by: Hulk Robot <hulkci@huawei.com>
      Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
      Acked-By: anton.ivanov@cambridgegreys.com
      Signed-off-by: Richard Weinberger <richard@nod.at>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b71f114670221c0175ea1c8484c02dff65d13a82
  Author: YiFei Zhu <zhuyifei1999@gmail.com>
  Date:   Tue Apr 20 00:56:10 2021 -0500

      um: Fix stack pointer alignment

      [ Upstream commit 558f9b2f94dbd2d5c5c8292aa13e081cc11ea7d9 ]

      GCC assumes that stack is aligned to 16-byte on call sites [1].
      Since GCC 8, GCC began using 16-byte aligned SSE instructions to
      implement assignments to structs on stack. When
      CC_OPTIMIZE_FOR_PERFORMANCE is enabled, this affects
      os-Linux/sigio.c, write_sigio_thread:

        struct pollfds *fds, tmp;
        tmp = current_poll;
```

```
    Note that struct pollfds is exactly 16 bytes in size.
    GCC 8+ generates assembly similar to:

      movdqa (%rdi),%xmm0
      movaps %xmm0,-0x50(%rbp)

    This is an issue, because movaps will #GP if -0x50(%rbp) is not
    aligned to 16 bytes [2], and how rbp gets assigned to is via glibc
    clone thread_start, then function prologue, going though execution
    trace similar to (showing only relevant instructions):

      sub    $0x10,%rsi
      mov    %rcx,0x8(%rsi)
      mov    %rdi,(%rsi)
      syscall
      pop    %rax
      pop    %rdi
      callq  *%rax
      push   %rbp
      mov    %rsp,%rbp

    The stack pointer always points to the topmost element on stack,
    rather then the space right above the topmost. On push, the
    pointer decrements first before writing to the memory pointed to
    by it. Therefore, there is no need to have the stack pointer
    pointer always point to valid memory unless the stack is poped;
    so the `- sizeof(void *)` in the code is unnecessary.

    On the other hand, glibc reserves the 16 bytes it needs on stack
    and pops itself, so by the call instruction the stack pointer
    is exactly the caller-supplied sp. It then push the 16 bytes of
    the return address and the saved stack pointer, so the base
    pointer will be 16-byte aligned if and only if the caller
    supplied sp is 16-byte aligned. Therefore, the caller must supply
    a 16-byte aligned pointer, which `stack + UM_KERN_PAGE_SIZE`
    already satisfies.

    On a side note, musl is unaffected by this issue because it forces
    16 byte alignment via `and $-16,%rsi` in its clone wrapper.
    Similarly, glibc i386 is also unaffected because it has
    `andl $0xfffffff0, %ecx`.

    To reproduce this bug, enable CONFIG_UML_RTC and
    CC_OPTIMIZE_FOR_PERFORMANCE. uml_rtc will call
    add_sigio_fd which will then cause write_sigio_thread to either go
    into segfault loop or panic with "Segfault with no mm".

    Similarly, signal stacks will be aligned by the host kernel upon
    signal delivery. `- sizeof(void *)` to sigaltstack is
    unconventional and extraneous.

    On a related note, initialization of longjmp buffers do require
    `- sizeof(void *)`. This is to account for the return address
    that would have been pushed to the stack at the call site.

    The reason for uml to respect 16-byte alignment, rather than
    telling GCC to assume 8-byte alignment like the host kernel since
    commit d9b0cde91c60 ("x86-64, gcc: Use
    -mpreferred-stack-boundary=3 if supported"), is because uml links
    against libc. There is no reason to assume libc is also compiled
    with that flag and assumes 8-byte alignment rather than 16-byte.

    [1] https://gcc.gnu.org/bugzilla/show_bug.cgi?id=40838
    [2] https://c9x.me/x86/html/file_module_x86_id_180.html
```

```
    Signed-off-by: YiFei Zhu <zhuyifei1999@gmail.com>
    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reviewed-by: Johannes Berg <johannes@sipsolutions.net>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a02357d7532b88e97329bd7786c7e72601109704
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Wed Jun 9 17:07:29 2021 -0400

    sunrpc: Avoid a KASAN slab-out-of-bounds bug in xdr_set_page_base()

    [ Upstream commit 6d1c0f3d28f98ea2736128ed3e46821496dc3a8c ]

    This seems to happen fairly easily during READ_PLUS testing on NFS v4.2.
    I found that we could end up accessing xdr->buf->pages[pgnr] with a pgnr
    greater than the number of pages in the array. So let's just return
    early if we're setting base to a point at the end of the page data and
    let xdr_set_tail_base() handle setting up the buffer pointers instead.

    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Fixes: 8d86e373b0ef ("SUNRPC: Clean up helpers xdr_set_iov() and xdr_set_page_base()")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7aec9f862411906f8c27071ba65a1e110ad7d2fd
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Fri Jun 11 13:40:55 2021 -0400

    NFSv4: Fix an Oops in pnfs_mark_request_commit() when doing O_DIRECT

    [ Upstream commit 3731d44bba8e0116b052b1b374476c5f6dd9a456 ]

    Fix an Oopsable condition in pnfs_mark_request_commit() when we're
    putting a set of writes on the commit list to reschedule them after a
    failed pNFS attempt.

    Fixes: 9c455a8c1e14 ("NFS/pNFS: Clean up pNFS commit operations")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b0bfac939030181177373f549398ba94c384713d
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Wed Jun 9 10:04:46 2021 -0400

    NFSv4: Initialise connection to the server in nfs4_alloc_client()

    [ Upstream commit dd99e9f98fbf423ff6d365b37a98e8879170f17c ]

    Set up the connection to the NFSv4 server in nfs4_alloc_client(), before
    we've added the struct nfs_client to the net-namespace's nfs_client_list
    so that a downed server won't cause other mounts to hang in the trunking
    detection code.

    Reported-by: Michael Wakabayashi <mwakabayashi@vmware.com>
    Fixes: 5c6e5b60aae4 ("NFS: Fix an Oops in the pNFS files and flexfiles connection setup
to the DS")
    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16187f5228e6171d8cd61e039ffeb88f46c3f7de
Author: Stephan Gerhold <stephan@gerhold.net>
Date:   Mon May 17 12:51:12 2021 +0200
```

```
    power: supply: rt5033_battery: Fix device tree enumeration

    [ Upstream commit f3076cd8d1d5fa64b5e1fa5affc045c2fc123baa ]

    The fuel gauge in the RT5033 PMIC has its own I2C bus and interrupt
    line. Therefore, it is not actually part of the RT5033 MFD and needs
    its own of_match_table to probe properly.

    Also, given that it's independent of the MFD, there is actually
    no need to make the Kconfig depend on MFD_RT5033. Although the driver
    uses the shared <linux/mfd/rt5033.h> header, there is no compile
    or runtime dependency on the RT5033 MFD driver.

    Cc: Beomho Seo <beomho.seo@samsung.com>
    Cc: Chanwoo Choi <cw00.choi@samsung.com>
    Fixes: b847dd96e659 ("power: rt5033_battery: Add RT5033 Fuel gauge device driver")
    Signed-off-by: Stephan Gerhold <stephan@gerhold.net>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7da9b60120b0883055c361308bad2309eb60d17e
Author: Krzysztof Wilczyński <kw@linux.com>
Date:   Thu Jun 3 00:01:12 2021 +0000

    PCI/sysfs: Fix dsm_label_utf16s_to_utf8s() buffer overrun

    [ Upstream commit bdcdaa13ad96f1a530711c29e6d4b8311eff767c ]

    "utf16s_to_utf8s(..., buf, PAGE_SIZE)" puts up to PAGE_SIZE bytes into
    "buf" and returns the number of bytes it actually put there.  If it wrote
    PAGE_SIZE bytes, the newline added by dsm_label_utf16s_to_utf8s() would
    overrun "buf".

    Reduce the size available for utf16s_to_utf8s() to use so there is always
    space for the newline.

    [bhelgaas: reorder patch in series, commit log]
    Fixes: 6058989bad05 ("PCI: Export ACPI _DSM provided firmware instance number and string
name to sysfs")
    Link: https://lore.kernel.org/r/20210603000112.703037-7-kw@linux.com
    Reported-by: Joe Perches <joe@perches.com>
    Signed-off-by: Krzysztof Wilczyński <kw@linux.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8fcb25a3ef232e615abfaceaf4a574776d5f1f8e
Author: Maximilian Luz <luzmaximilian@gmail.com>
Date:   Tue May 11 11:24:21 2021 +0200

    power: supply: surface-charger: Fix type of integer variable

    [ Upstream commit 601423bc0c06467d019cf2a446962a5bf1b5e330 ]

    The ac->state field is __le32, not u32. So change the variable we're
    temporarily storing it in to __le32 as well.

    Reported-by: kernel test robot <lkp@intel.com>
    Fixes: e61ffb344591 ("power: supply: Add AC driver for Surface Aggregator Module")
    Signed-off-by: Maximilian Luz <luzmaximilian@gmail.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 10d79d3bf5c45c8a82a88d3076206628856d59af
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Thu May 6 22:46:01 2021 +0200
```

```
    remoteproc: k3-r5: Fix an error message

    [ Upstream commit 34c4da6d5dfba48f49f891ebd75bb55999f0c538 ]

    'ret' is known to be 0 here.
    Reorder the code so that the expected error code is printed.

    Acked-by: Suman Anna <s-anna@ti.com>
    Fixes: 6dedbd1d5443 ("remoteproc: k3-r5: Add a remoteproc driver for R5F subsystem")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Link:
https://lore.kernel.org/r/d6e29d903b48957bf59c67229d54b0fc215e31ae.1620333870.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f02de04e14cf2ed7172300dfb5fe4d996b3adfac
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Wed Apr 21 16:00:40 2021 +0200

    remoteproc: stm32: fix phys_addr_t format string

    [ Upstream commit 3e25e407a1c93b53a87a7743ea0cd4703d3985b7 ]

    A phys_addr_t may be wider than an int or pointer:

    drivers/remoteproc/stm32_rproc.c: In function 'stm32_rproc_da_to_pa':
    drivers/remoteproc/stm32_rproc.c:583:30: error: format '%x' expects argument of type
'unsigned int', but argument 5 has type 'phys_addr_t' {aka 'long long unsigned int'} [-
Werror=format=]
      583 |                dev_dbg(dev, "da %llx to pa %#x\n", da, *pa);

    Print it by reference using the special %pap format string.

    Reviewed-by: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
    Fixes: 8a471396d21c ("remoteproc: stm32: Move resource table setup to rproc_ops")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20210421140053.3727528-1-arnd@kernel.org
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ca41d7ba671e868d00cc7c2aee020dd54a75e2b
Author: Chao Yu <chao@kernel.org>
Date:   Tue May 18 17:54:58 2021 +0800

    f2fs: compress: fix to disallow temp extension

    [ Upstream commit 4a67d9b07ac8dce7f1034e0d887f2f4ee00fe118 ]

    This patch restricts to configure compress extension as format of:

     [filename + '.' + extension]

    rather than:

     [filename + '.' + extension + (optional: '.' + temp extension)]

    in order to avoid to enable compression incorrectly:

    1. compress_extension=so
    2. touch file.soa
    3. touch file.so.tmp

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

```
    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3988657d3703df7051b46e57250306ab82aed6e
Author: Chao Yu <chao@kernel.org>
Date:   Tue May 18 09:57:54 2021 +0800

    f2fs: add MODULE_SOFTDEP to ensure crc32 is included in the initramfs

    [ Upstream commit 0dd571785d61528d62cdd8aa49d76bc6085152fe ]

    As marcosfrm reported in bugzilla:

    https://bugzilla.kernel.org/show_bug.cgi?id=213089

    Initramfs generators rely on "pre" softdeps (and "depends") to include
    additional required modules.

    F2FS does not declare "pre: crc32" softdep. Then every generator (dracut,
    mkinitcpio...) has to maintain a hardcoded list for this purpose.

    Hence let's use MODULE_SOFTDEP("pre: crc32") in f2fs code.

    Fixes: 43b6573bac95 ("f2fs: use cryptoapi crc32 functions")
    Reported-by: marcosfrm <marcosfrm@gmail.com>
    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fa815e11816b1fd1524829c693429c2e53290558
Author: Jon Mediero <jmdr@disroot.org>
Date:   Thu May 20 14:23:26 2021 +0200

    module: correctly exit module_kallsyms_on_each_symbol when fn() != 0

    [ Upstream commit 2c0f0f3639562d6e38ee9705303c6457c4936eac ]

    Commit 013c1667cf78 ("kallsyms: refactor
    {,module_}kallsyms_on_each_symbol") replaced the return inside the
    nested loop with a break, changing the semantics of the function: the
    break only exits the innermost loop, so the code continues iterating the
    symbols of the next module instead of exiting.

    Fixes: 013c1667cf78 ("kallsyms: refactor {,module_}kallsyms_on_each_symbol")
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Reviewed-by: Miroslav Benes <mbenes@suse.cz>
    Signed-off-by: Jon Mediero <jmdr@disroot.org>
    Signed-off-by: Jessica Yu <jeyu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7c2089413c675101652860599c69e8424647c10b
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Mon Apr 26 17:03:50 2021 +0200

    pwm: visconti: Fix and simplify period calculation

    [ Upstream commit 937efa29e70f7f8424b74631375dcb35d82a4614 ]

    With the original code a request for period = 65536000 ns and period =
    32768000 ns yields the same register settings (which results in 32768000
    ns) because the value for pwmc0 was miscalculated.

    Also simplify using that fls(0) is 0.
```

```
    Fixes: 721b595744f1 ("pwm: visconti: Add Toshiba Visconti SoC PWM support")
    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Acked-by: Nobuhiro Iwamatsu <nobuhiro1.iwamatsu@toshiba.co.jp>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee24633a50483e30f16a9edff11f25038507b266
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Thu May 6 21:09:48 2021 +0200

    cpufreq: scmi: Fix an error message

    [ Upstream commit b791c7f94680ba9b60b0c0786b1d0eb4393053d6 ]

    'ret' is known to be 0 here.
    The last error code is stored in 'nr_opp', so use it in the error message.

    Fixes: 71a37cd6a59d ("scmi-cpufreq: Remove deferred probe")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Reviewed-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit afb04d0b5543a5bf8e157b9119fbfc52606f4c11
Author: Chang S. Bae <chang.seok.bae@intel.com>
Date:   Tue May 18 13:03:19 2021 -0700

    x86/signal: Detect and prevent an alternate signal stack overflow

    [ Upstream commit 2beb4a53fc3f1081cedc1c1a198c7f56cc4fc60c ]

    The kernel pushes context on to the userspace stack to prepare for the
    user's signal handler. When the user has supplied an alternate signal
    stack, via sigaltstack(2), it is easy for the kernel to verify that the
    stack size is sufficient for the current hardware context.

    Check if writing the hardware context to the alternate stack will exceed
    it's size. If yes, then instead of corrupting user-data and proceeding with
    the original signal handler, an immediate SIGSEGV signal is delivered.

    Refactor the stack pointer check code from on_sig_stack() and use the new
    helper.

    While the kernel allows new source code to discover and use a sufficient
    alternate signal stack size, this check is still necessary to protect
    binaries with insufficient alternate signal stack size from data
    corruption.

    Fixes: c2bc11f10a39 ("x86, AVX-512: Enable AVX-512 States Context Switch")
    Reported-by: Florian Weimer <fweimer@redhat.com>
    Suggested-by: Jann Horn <jannh@google.com>
    Suggested-by: Andy Lutomirski <luto@kernel.org>
    Signed-off-by: Chang S. Bae <chang.seok.bae@intel.com>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Reviewed-by: Len Brown <len.brown@intel.com>
    Acked-by: Thomas Gleixner <tglx@linutronix.de>
    Link: https://lkml.kernel.org/r/20210518200320.17239-6-chang.seok.bae@intel.com
    Link: https://bugzilla.kernel.org/show_bug.cgi?id=153531
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e0648a3cb639e760b079d56c685a837dadee7a7
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Fri May 14 15:55:17 2021 -0400

    NFSD: Fix TP_printk() format specifier in nfsd_clid_class
```

```
    [ Upstream commit a948b1142cae66785521a389cab2cce74069b547 ]

    Since commit 9a6944fee68e ("tracing: Add a verifier to check string
    pointers for trace events"), which was merged in v5.13-rc1,
    TP_printk() no longer tacitly supports the "%.*s" format specifier.

    These are low value tracepoints, so just remove them.

    Reported-by: David Wysochanski <dwysocha@redhat.com>
    Fixes: dd5e3fbc1f47 ("NFSD: Add tracepoints to the NFSD state management code")
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Cc: Steven Rostedt <rostedt@goodmis.org>
    Signed-off-by: J. Bruce Fields <bfields@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1f1f6603155ae7dba7f68bbf2621566601f0fde
Author: Chao Yu <chao@kernel.org>
Date:   Tue May 11 18:17:34 2021 +0800

    f2fs: atgc: fix to set default age threshold

    [ Upstream commit 89e53ff1651a61cf2abef9356e2f60d0086215be ]

    Default age threshold value is missed to set, fix it.

    Fixes: 093749e296e2 ("f2fs: support age threshold based garbage collection")
    Reported-by: Sahitya Tummala <stummala@codeaurora.org>
    Signed-off-by: Chao Yu <yuchao0@huawei.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f0fd04efc26dbbda1ba0dbf9f76facc0d9a5c1a
Author: Maximilian Luz <luzmaximilian@gmail.com>
Date:   Tue May 4 20:48:13 2021 +0200

    power: supply: surface_battery: Fix battery event handling

    [ Upstream commit e633f33d2669cb54db2846f9cde08662d254dbd3 ]

    The battery subsystem of the Surface Aggregator Module EC requires us to
    register the battery notifier with instance ID 0. However, battery
    events are actually sent with the instance ID corresponding to the
    device, which is nonzero. Thus, the strict-matching approach doesn't
    work here and will discard events that the driver is expected to handle.

    To fix this we have to fall back on notifier matching by target-category
    only and have to manually check the instance ID in the notifier
    callback.

    Fixes: 167f77f7d0b3 ("power: supply: Add battery driver for Surface Aggregator Module")
    Signed-off-by: Maximilian Luz <luzmaximilian@gmail.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 73a14f65a67c093771a3180f523f8b11e729417f
Author: Chunguang Xu <brookxu@tencent.com>
Date:   Tue Jul 6 05:47:26 2021 +0800

    block: fix the problem of io_ticks becoming smaller

    [ Upstream commit d80c228d44640f0b47b57a2ca4afa26ef87e16b0 ]

    On the IO submission path, blk_account_io_start() may interrupt
    the system interruption. When the interruption returns, the value
```

```
    of part->stamp may have been updated by other cores, so the time
    value collected before the interruption may be less than part->
    stamp. So when this happens, we should do nothing to make io_ticks
    more accurate? For kernels less than 5.0, this may cause io_ticks
    to become smaller, which in turn may cause abnormal ioutil values.

    Signed-off-by: Chunguang Xu <brookxu@tencent.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/1625521646-1069-1-git-send-email-brookxu.cn@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 21a06a244d2576f93cbc9ce9bf95814c2810c36a
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Tue May 25 20:56:22 2021 +0800

    virtio_console: Assure used length from device is limited

    [ Upstream commit d00d8da5869a2608e97cfede094dfc5e11462a46 ]

    The buf->len might come from an untrusted device. This
    ensures the value would not exceed the size of the buffer
    to avoid data corruption or loss.

    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Acked-by: Jason Wang <jasowang@redhat.com>
    Link: https://lore.kernel.org/r/20210525125622.1203-1-xieyongji@bytedance.com
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffbfea49cc39adba0a26ab2a172c0c86a7130a35
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Mon May 17 16:45:16 2021 +0800

    virtio_net: Fix error handling in virtnet_restore()

    [ Upstream commit 3f2869cace829fb4b80fc53b3ddaa7f4ba9acbf1 ]

    Do some cleanups in virtnet_restore() when virtnet_cpu_notif_add() failed.

    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Link: https://lore.kernel.org/r/20210517084516.332-1-xieyongji@bytedance.com
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29a2f4a3214aa14d61cc9737c9f886dae9dbb710
Author: Xie Yongji <xieyongji@bytedance.com>
Date:   Mon May 17 16:43:32 2021 +0800

    virtio-blk: Fix memory leak among suspend/resume procedure

    [ Upstream commit b71ba22e7c6c6b279c66f53ee7818709774efa1f ]

    The vblk->vqs should be freed before we call init_vqs()
    in virtblk_restore().

    Signed-off-by: Xie Yongji <xieyongji@bytedance.com>
    Link: https://lore.kernel.org/r/20210517084332.280-1-xieyongji@bytedance.com
    Acked-by: Jason Wang <jasowang@redhat.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1c8c1adae69f19e8e699ae9f34c737e4fe374999
Author: Ye Bin <yebin10@huawei.com>
```

```
Date:   Tue Jun 15 17:05:37 2021 +0800

    ext4: fix WARN_ON_ONCE(!buffer_uptodate) after an error writing the superblock

    [ Upstream commit 558d6450c7755aa005d89021204b6cdcae5e848f ]

    If a writeback of the superblock fails with an I/O error, the buffer
    is marked not uptodate.  However, this can cause a WARN_ON to trigger
    when we attempt to write superblock a second time.  (Which might
    succeed this time, for cerrtain types of block devices such as iSCSI
    devices over a flaky network.)

    Try to detect this case in flush_stashed_error_work(), and also change
    __ext4_handle_dirty_metadata() so we always set the uptodate flag, not
    just in the nojournal case.

    Before this commit, this problem can be repliciated via:

    1. dmsetup  create dust1 --table  '0 2097152 dust /dev/sdc 0 4096'
    2. mount  /dev/mapper/dust1  /home/test
    3. dmsetup message dust1 0 addbadblock 0 10
    4. cd /home/test
    5. echo "XXXXXXX" > t

    After a few seconds, we got following warning:

    [   80.654487] end_buffer_async_write: bh=0xffff88842f18bdd0
    [   80.656134] Buffer I/O error on dev dm-0, logical block 0, lost async page write
    [   85.774450] EXT4-fs error (device dm-0): ext4_check_bdev_write_error:193: comm
kworker/u16:8: Error while async write back metadata
    [   91.415513] mark_buffer_dirty: bh=0xffff88842f18bdd0
    [   91.417038] ------------[ cut here ]------------
    [   91.418450] WARNING: CPU: 1 PID: 1944 at fs/buffer.c:1092
mark_buffer_dirty.cold+0x1c/0x5e
    [   91.440322] Call Trace:
    [   91.440652]  __jbd2_journal_temp_unlink_buffer+0x135/0x220
    [   91.441354]  __jbd2_journal_unfile_buffer+0x24/0x90
    [   91.441981]  __jbd2_journal_refile_buffer+0x134/0x1d0
    [   91.442628]  jbd2_journal_commit_transaction+0x249a/0x3240
    [   91.443336]  ? put_prev_entity+0x2a/0x200
    [   91.443856]  ? kjournald2+0x12e/0x510
    [   91.444324]  kjournald2+0x12e/0x510
    [   91.444773]  ? woken_wake_function+0x30/0x30
    [   91.445326]  kthread+0x150/0x1b0
    [   91.445739]  ? commit_timeout+0x20/0x20
    [   91.446258]  ? kthread_flush_worker+0xb0/0xb0
    [   91.446818]  ret_from_fork+0x1f/0x30
    [   91.447293] ---[ end trace 66f0b6bf3d1abade ]---

    Signed-off-by: Ye Bin <yebin10@huawei.com>
    Link: https://lore.kernel.org/r/20210615090537.3423231-1-yebin10@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b881ba731274645976d569cfeb4dd6a8a8cde5bd
Author: Javier Martinez Canillas <javierm@redhat.com>
Date:   Tue Jun 8 10:04:09 2021 +0200

    PCI: rockchip: Register IRQ handlers after device and data are ready

    [ Upstream commit 3cf5f7ab230e2b886e493c7a8449ed50e29d2b98 ]

    An IRQ handler may be called at any time after it is registered, so
    anything it relies on must be ready before registration.
```

```
    rockchip_pcie_subsys_irq_handler() and rockchip_pcie_client_irq_handler()
    read registers in the PCIe controller, but we registered them before
    turning on clocks to the controller.  If either is called before the clocks
    are turned on, the register reads fail and the machine hangs.

    Similarly, rockchip_pcie_legacy_int_handler() uses rockchip->irq_domain,
    but we installed it before initializing irq_domain.

    Register IRQ handlers after their data structures are initialized and
    clocks are enabled.

    Found by enabling CONFIG_DEBUG_SHIRQ, which calls the IRQ handler when it
    is being unregistered.  An error during the probe path might cause this
    unregistration and IRQ handler execution before the device or data
    structure init has finished.

    [bhelgaas: commit log]
    Link: https://lore.kernel.org/r/20210608080409.1729276-1-javierm@redhat.com
    Reported-by: Peter Robinson <pbrobinson@gmail.com>
    Tested-by: Peter Robinson <pbrobinson@gmail.com>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Shawn Lin <shawn.lin@rock-chips.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfac32923d1de4e2e14004b8a7d6e3caf4fdef72
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Wed Jun 30 17:23:16 2021 +0200

    ACPI: video: Add quirk for the Dell Vostro 3350

    [ Upstream commit 9249c32ec9197e8d34fe5179c9e31668a205db04 ]

    The Dell Vostro 3350 ACPI video-bus device reports spurious
    ACPI_VIDEO_NOTIFY_CYCLE events resulting in spurious KEY_SWITCHVIDEOMODE
    events being reported to userspace (and causing trouble there).

    Add a quirk setting the report_key_events mask to
    REPORT_BRIGHTNESS_KEY_EVENTS so that the ACPI_VIDEO_NOTIFY_CYCLE
    events will be ignored, while still reporting brightness up/down
    hotkey-presses to userspace normally.

    BugLink: https://bugzilla.redhat.com/show_bug.cgi?id=1911763
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16752e05c77a233f23a9fe8acebfb81889d7c116
Author: Liguang Zhang <zhangliguang@linux.alibaba.com>
Date:   Tue Jun 29 19:27:48 2021 +0800

    ACPI: AMBA: Fix resource name in /proc/iomem

    [ Upstream commit 7718629432676b5ebd9a32940782fe297a0abf8d ]

    In function amba_handler_attach(), dev->res.name is initialized by
    amba_device_alloc. But when address_found is false, dev->res.name is
    assigned to null value, which leads to wrong resource name display in
    /proc/iomem, "<BAD>" is seen for those resources.

    Signed-off-by: Liguang Zhang <zhangliguang@linux.alibaba.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit fd313d643f07810334091d1535c3cd861b723bcc
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Thu Jun 17 11:51:41 2021 +0200

    pwm: tegra: Don't modify HW state in .remove callback

    [ Upstream commit 86f7fa71cd830d18d7ebcaf719dffd5ddfe1acdd ]

    A consumer is expected to disable a PWM before calling pwm_put(). And if
    they didn't there is hopefully a good reason (or the consumer needs
    fixing). Also if disabling an enabled PWM was the right thing to do,
    this should better be done in the framework instead of in each low level
    driver.

    So drop the hardware modification from the .remove() callback.

    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a1d3f7407f78fa7d6d801e236dd26f73bf03e0a4
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 11:57:17 2021 +0800

    pwm: img: Fix PM reference leak in img_pwm_enable()

    [ Upstream commit fde25294dfd8e36e4e30b693c27a86232864002a ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 18127c30f61955382abd9e8317b5eb438278e028
Author: Philip Yang <Philip.Yang@amd.com>
Date:   Mon Jun 21 18:51:26 2021 -0400

    drm/amdkfd: fix sysfs kobj leak

    [ Upstream commit dcdb4d904b4bd3078fe8d4d24b1658560d6078ef ]

    3 cases of kobj leak, which causes memory leak:

    kobj_type must have release() method to free memory from release
    callback. Don't need NULL default_attrs to init kobj.

    sysfs files created under kobj_status should be removed with kobj_status
    as parent kobject.

    Remove queue sysfs files when releasing queue from process MMU notifier
    release callback.

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

commit 0b909924a4aeeb6fa9b0827bbc855d570ef99770
Author: Evan Quan <evan.quan@amd.com>
Date:   Tue May 25 11:43:38 2021 +0800
```

```
    drm/amdgpu: fix Navi1x tcp power gating hang when issuing lightweight invalidaiton

    [ Upstream commit 9c26ddb1c5b6e30c6bca48b8ad9205d96efe93d0 ]

    Fix TCP hang when a lightweight invalidation happens on Navi1x.

    Signed-off-by: Evan Quan <evan.quan@amd.com>
    Reviewed-by: Lijo Lazar <lijo.lazar@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 813ff24253242af9823658147ceffe5124b15fc1
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Jun 17 21:27:02 2021 +0200

    power: supply: axp288_fuel_gauge: Make "T3 MRD" no_battery_list DMI entry more generic

    [ Upstream commit 3a06b912a5ce494d7b7300b12719c562be7b566f ]

    It turns out that the "T3 MRD" DMI_BOARD_NAME value is used in a lot of
    different Cherry Trail x5-z8300 / x5-z8350 based Mini-PC / HDMI-stick
    models from Ace PC / Meegopad / MinisForum / Wintel (and likely also
    other vendors).

    Most of the other DMI strings on these boxes unfortunately contain various
    generic values like "Default string" or "$(DEFAULT_STRING)", so we cannot
    match on them. These devices do have their chassis-type correctly set to a
    value of "3" (desktop) which is a pleasant surprise, so also match on that.

    This should avoid the quirk accidentally also getting applied to laptops /
    tablets (which do actually have a battery). Although in my quite large
    database of Bay and Cherry Trail based devices DMIdecode dumps I don't
    have any laptops / tables with a board-name of "T3 MRD", so this should
    not be an issue.

    BugLink: https://askubuntu.com/questions/1206714/how-can-a-mini-pc-be-stopped-from-
being-detected-as-a-laptop-with-a-battery/
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2898347605790d76915ff745e795064746cc1007
Author: Zou Wei <zou_wei@huawei.com>
Date:   Sat Jun 5 09:21:41 2021 +0800

    power: supply: ab8500: add missing MODULE_DEVICE_TABLE

    [ Upstream commit dfe52db13ab8d24857a9840ec7ca75eef800c26c ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 81cd77390071d8fe081de1382ef7c7bda82e74b1
Author: Zou Wei <zou_wei@huawei.com>
Date:   Sat Jun 5 09:21:54 2021 +0800

    power: supply: charger-manager: add missing MODULE_DEVICE_TABLE
```

```
    [ Upstream commit 073b5d5b1f9cc94a3eea25279fbafee3f4f5f097 ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a01f95e2f881c163d0bd9b4efe91a278e045763b
Author: Zou Wei <zou_wei@huawei.com>
Date:   Sat Jun 5 09:21:23 2021 +0800

    power: reset: regulator-poweroff: add missing MODULE_DEVICE_TABLE

    [ Upstream commit 4465b3a621e761d82d1a92e3fda88c5d33c804b8 ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c11910b2c2cff64847d438c83d88e6e3c282c91b
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Tue May 11 23:41:10 2021 -0400

    NFS: nfs_find_open_context() may only select open files

    [ Upstream commit e97bc66377bca097e1f3349ca18ca17f202ff659 ]

    If a file has already been closed, then it should not be selected to
    support further I/O.

    Signed-off-by: Trond Myklebust <trond.myklebust@hammerspace.com>
    [Trond: Fix an invalid pointer deref reported by Colin Ian King]
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a54caa8540fac6ac3297b4eb5b5832ea7fd7b56
Author: Jing Xiangfeng <jingxiangfeng@huawei.com>
Date:   Tue Jun 29 19:59:56 2021 +0800

    drm/gma500: Add the missed drm_gem_object_put() in psb_user_framebuffer_create()

    [ Upstream commit cd8f318fbd266b127ffc93cc4c1eaf9a5196fafb ]

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

    Signed-off-by: Jing Xiangfeng <jingxiangfeng@huawei.com>
    Signed-off-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210629115956.15160-1-
jingxiangfeng@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e43aa62408141f366bec1a93d590bd15340310f8
Author: Jeff Layton <jlayton@kernel.org>
Date:   Tue May 4 10:08:30 2021 -0400

    ceph: remove bogus checks and WARN_ONs from ceph_set_page_dirty
```

```
    [ Upstream commit 22d41cdcd3cfd467a4af074165357fcbea1c37f5 ]

    The checks for page->mapping are odd, as set_page_dirty is an
    address_space operation, and I don't see where it would be called on a
    non-pagecache page.

    The warning about the page lock also seems bogus.  The comment over
    set_page_dirty() says that it can be called without the page lock in
    some rare cases. I don't think we want to warn if that's the case.

    Reported-by: Matthew Wilcox <willy@infradead.org>
    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Ilya Dryomov <idryomov@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4f8e35e9483a894b1c256b51c8dc32cb56a1df9c
Author: Mike Marshall <hubcap@omnibond.com>
Date:   Tue May 18 08:09:13 2021 -0400

    orangefs: fix orangefs df output.

    [ Upstream commit 0fdec1b3c9fbb5e856a40db5993c9eaf91c74a83 ]

    Orangefs df output is whacky. Walt Ligon suggested this might fix it.
    It seems way more in line with reality now...

    Signed-off-by: Mike Marshall <hubcap@omnibond.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cdc00734394b14832589b76eeb0440bc8bc6bc10
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Fri Jun 25 15:49:31 2021 -0400

    NFSv4: Fix handling of non-atomic change attrbute updates

    [ Upstream commit 20cf7d4ea4ad7d9830b01ff7444f6ac64a727a23 ]

    If the change attribute update is declared to be non-atomic by the
    server, or our cached value does not match the server's value before the
    operation was performed, then we should declare the inode cache invalid.

    On the other hand, if the change to the directory raced with a lookup or
    getattr which already updated the change attribute, then optimise away
    the revalidation.

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

commit 357d8812007273301ac14de2c3fbc2eabf515597
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Sat Jun 26 09:23:51 2021 -0400

    NFS: Fix up inode attribute revalidation timeouts

    [ Upstream commit 213bb58475b57786e4336bc8bfd5029e16257c49 ]

    The inode is considered revalidated when we've checked the value of the
    change attribute against our cached value since that suffices to
    establish whether or not the other cached values are valid.

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

commit 6e56ae0a890170a369dece7efd8c0fa97974a186
```

```
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 12:07:02 2021 +0800

    PCI: tegra: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 7bf475a4614a9722b9b989e53184a02596cf16d1 ]

    Add missing MODULE_DEVICE_TABLE definition so we generate correct modalias
    for automatic loading of this driver when it is built as a module.

    Link: https://lore.kernel.org/r/1620792422-16535-1-git-send-email-zou_wei@huawei.com
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Vidya Sagar <vidyas@nvidia.com>
    Acked-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7693312a80445911b298c9a71eae047fb2750c2f
Author: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
Date:   Tue Apr 20 11:19:22 2021 +0200

    remoteproc: stm32: fix mbox_send_message call

    [ Upstream commit 51c4b4e212269a8634dee2000182cfca7f11575b ]

    mbox_send_message is called by passing a local dummy message or
    a function parameter. As the message is queued, it is dereferenced.
    This works because the message field is not used by the stm32 ipcc
    driver, but it is not clean.

    Fix by passing a constant string in all cases.

    The associated comments are removed because rproc should not have to
    deal with the behavior of the mailbox frame.

    Reported-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
    Link: https://lore.kernel.org/r/20210420091922.29429-1-arnaud.pouliquen@foss.st.com
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a9f2ab385cc6942d02791eb46c15f79cb46253ac
Author: Siddharth Gupta <sidgup@codeaurora.org>
Date:   Mon Jun 14 19:21:10 2021 -0700

    remoteproc: core: Fix cdev remove and rproc del

    [ Upstream commit 930eec0be20c93a53160c74005a1485a230e6911 ]

    The rproc_char_device_remove() call currently unmaps the cdev
    region instead of simply deleting the cdev that was added as a
    part of the rproc_char_device_add() call. This change fixes that
    behaviour, and also fixes the order in which device_del() and
    cdev_del() need to be called.

    Signed-off-by: Siddharth Gupta <sidgup@codeaurora.org>
    Link: https://lore.kernel.org/r/1623723671-5517-4-git-send-email-sidgup@codeaurora.org
    Signed-off-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbea0e7897708173464395a6f9478f1b17f84f55
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Jun 23 14:02:30 2021 +0200
```

```
    x86/fpu: Return proper error codes from user access functions

    [ Upstream commit aee8c67a4faa40a8df4e79316dbfc92d123989c1 ]

    When *RSTOR from user memory raises an exception, there is no way to
    differentiate them. That's bad because it forces the slow path even when
    the failure was not a fault. If the operation raised eg. #GP then going
    through the slow path is pointless.

    Use _ASM_EXTABLE_FAULT() which stores the trap number and let the exception
    fixup return the negated trap number as error.

    This allows to separate the fast path and let it handle faults directly and
    avoid the slow path for all other exceptions.

    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Signed-off-by: Borislav Petkov <bp@suse.de>
    Link: https://lkml.kernel.org/r/20210623121457.601480369@linutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8a51030e3be01b2eb721718e3c22d571b77e505
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 15:11:31 2021 +0800

    PCI: mediatek-gen3: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 3a2e476dc5d02af3422143b07d8db1eced475314 ]

    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.

    Link: https://lore.kernel.org/r/1620717091-108691-1-git-send-email-zou_wei@huawei.com
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Krzysztof Wilczyński <kw@linux.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c09a4ad6251f9354df0915ac5b8f22c7b9d6bc11
Author: Amir Goldstein <amir73il@gmail.com>
Date:   Mon Jun 21 14:03:53 2021 +0300

    fuse: fix illegal access to inode with reused nodeid

    [ Upstream commit 15db16837a35d8007cb8563358787412213db25e ]

    Server responds to LOOKUP and other ops (READDIRPLUS/CREATE/MKNOD/...)
    with ourarg containing nodeid and generation.

    If a fuse inode is found in inode cache with the same nodeid but different
    generation, the existing fuse inode should be unhashed and marked "bad" and
    a new inode with the new generation should be hashed instead.

    This can happen, for example, with passhrough fuse filesystem that returns
    the real filesystem ino/generation on lookup and where real inode numbers
    can get recycled due to real files being unlinked not via the fuse
    passthrough filesystem.

    With current code, this situation will not be detected and an old fuse
    dentry that used to point to an older generation real inode, can be used to
    access a completely new inode, which should be accessed only via the new
    dentry.
```

```
    Note that because the FORGET message carries the nodeid w/o generation, the
    server should wait to get FORGET counts for the nlookup counts of the old
    and reused inodes combined, before it can free the resources associated to
    that nodeid.

    Signed-off-by: Amir Goldstein <amir73il@gmail.com>
    Signed-off-by: Miklos Szeredi <mszeredi@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6ba041fc3c441e0cf4762f4baf0166e8d34f6802
Author: Greg Kurz <groug@kaod.org>
Date:   Thu May 20 17:46:54 2021 +0200

    virtiofs: propagate sync() to file server

    [ Upstream commit 2d82ab251ef0f6e7716279b04e9b5a01a86ca530 ]

    Even if POSIX doesn't mandate it, linux users legitimately expect sync() to
    flush all data and metadata to physical storage when it is located on the
    same system.  This isn't happening with virtiofs though: sync() inside the
    guest returns right away even though data still needs to be flushed from
    the host page cache.

    This is easily demonstrated by doing the following in the guest:

    $ dd if=/dev/zero of=/mnt/foo bs=1M count=5K ; strace -T -e sync sync
    5120+0 records in
    5120+0 records out
    5368709120 bytes (5.4 GB, 5.0 GiB) copied, 5.22224 s, 1.0 GB/s
    sync()                                  = 0 <0.024068>

    and start the following in the host when the 'dd' command completes
    in the guest:

    $ strace -T -e fsync /usr/bin/sync virtiofs/foo
    fsync(3)                                = 0 <10.371640>

    There are no good reasons not to honor the expected behavior of sync()
    actually: it gives an unrealistic impression that virtiofs is super fast
    and that data has safely landed on HW, which isn't the case obviously.

    Implement a ->sync_fs() superblock operation that sends a new FUSE_SYNCFS
    request type for this purpose.  Provision a 64-bit placeholder for possible
    future extensions.  Since the file server cannot handle the wait == 0 case,
    we skip it to avoid a gratuitous roundtrip.  Note that this is
    per-superblock: a FUSE_SYNCFS is send for the root mount and for each
    submount.

    Like with FUSE_FSYNC and FUSE_FSYNCDIR, lack of support for FUSE_SYNCFS in
    the file server is treated as permanent success.  This ensures
    compatibility with older file servers: the client will get the current
    behavior of sync() not being propagated to the file server.

    Note that such an operation allows the file server to DoS sync().  Since a
    typical FUSE file server is an untrusted piece of software running in
    userspace, this is disabled by default.  Only enable it with virtiofs for
    now since virtiofsd is supposedly trusted by the guest kernel.

    Reported-by: Robert Krawitz <rlk@redhat.com>
    Signed-off-by: Greg Kurz <groug@kaod.org>
    Signed-off-by: Miklos Szeredi <mszeredi@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e65819a006ec8a8df2f8639dc26ef0cfaa95ae7
Author: Jan Kiszka <jan.kiszka@siemens.com>
```

```
Date:   Sun May 30 13:24:23 2021 +0200

    watchdog: iTCO_wdt: Account for rebooting on second timeout

    [ Upstream commit cb011044e34c293e139570ce5c01aed66a34345c ]

    This was already attempted to fix via 1fccb73011ea: If the BIOS did not
    enable TCO SMIs, the timer definitely needs to trigger twice in order to
    cause a reboot. If TCO SMIs are on, as well as SMIs in general, we can
    continue to assume that the BIOS will perform a reboot on the first
    timeout.

    QEMU with its ICH9 and related BIOS falls into the former category,
    currently taking twice the configured timeout in order to reboot the
    machine. For iTCO version that fall under turn_SMI_watchdog_clear_off,
    this is also true and was currently only addressed for v1, irrespective
    of the turn_SMI_watchdog_clear_off value.

    Signed-off-by: Jan Kiszka <jan.kiszka@siemens.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/0b8bb307-d08b-41b5-696c-305cdac6789c@siemens.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cfb7e2f23aa063b6a29df2badc942904fa47a0df
Author: Stefan Eichenberger <eichest@gmail.com>
Date:   Tue Apr 6 14:12:47 2021 +0200

    watchdog: imx_sc_wdt: fix pretimeout

    [ Upstream commit 854478a381078ee86ae2a7908a934b1ded399130 ]

    If the WDIOF_PRETIMEOUT flag is not set when registering the device the
    driver will not show the sysfs entries or register the default governor.
    By moving the registering after the decision whether pretimeout is
    supported this gets fixed.

    Signed-off-by: Stefan Eichenberger <eichest@gmail.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Reviewed-by: Dong Aisheng <aisheng.dong@nxp.com>
    Link: https://lore.kernel.org/r/20210519080311.142928-1-eichest@gmail.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a053c4d716898a53c2e31c574a70ea0c37044a3
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 14:57:56 2021 +0800

    watchdog: Fix possible use-after-free by calling del_timer_sync()

    [ Upstream commit d0212f095ab56672f6f36aabc605bda205e1e0bf ]

    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.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
```

```
    Acked-by: Vladimir Zapolskiy <vz@mleia.com>
    Link: https://lore.kernel.org/r/1620802676-19701-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f0feab82f6a0323f54d85e8b512a2be64f83648a
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 15:04:51 2021 +0800

    watchdog: sc520_wdt: Fix possible use-after-free in wdt_turnoff()

    [ Upstream commit 90b7c141132244e8e49a34a4c1e445cce33e07f4 ]

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

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/1620716691-108460-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8adbbe6c86bb13e14f8a19e036ae5f4f5661fd90
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 15:01:35 2021 +0800

    watchdog: Fix possible use-after-free in wdt_startup()

    [ Upstream commit c08a6b31e4917034f0ed0cb457c3bb209576f542 ]

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

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/1620716495-108352-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Wim Van Sebroeck <wim@linux-watchdog.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a037ebbe72a4f98495b193112e2b2000e5e09eb5
Author: Russell King <rmk+kernel@armlinux.org.uk>
Date:   Thu May 13 15:18:27 2021 +0100

    PCI: Dynamically map ECAM regions

    [ Upstream commit 8fe55ef23387ce3c7488375b1fd539420d7654bb ]

    Attempting to boot 32-bit ARM kernels under QEMU's 3.x virt models fails
    when we have more than 512M of RAM in the model as we run out of vmalloc
    space for the PCI ECAM regions. This failure will be silent when running
```

```
    libvirt, as the console in that situation is a PCI device.

    In this configuration, the kernel maps the whole ECAM, which QEMU sets up
    for 256 buses, even when maybe only seven buses are in use.  Each bus uses
    1M of ECAM space, and ioremap() adds an additional guard page between
    allocations. The kernel vmap allocator will align these regions to 512K,
    resulting in each mapping eating 1.5M of vmalloc space. This means we need
    384M of vmalloc space just to map all of these, which is very wasteful of
    resources.

    Fix this by only mapping the ECAM for buses we are going to be using.  In
    my setups, this is around seven buses in most guests, which is 10.5M of
    vmalloc space - way smaller than the 384M that would otherwise be required.
    This also means that the kernel can boot without forcing extra RAM into
    highmem with the vmalloc= argument, or decreasing the virtual RAM available
    to the guest.

    Suggested-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/E1lhCAV-0002yb-50@rmk-PC.armlinux.org.uk
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3cf262e8c6980ac300e01624f4dc9a9f85868b91
Author: Lukas Wunner <lukas@wunner.de>
Date:   Sat May 1 10:29:00 2021 +0200

    PCI: pciehp: Ignore Link Down/Up caused by DPC

    [ Upstream commit a97396c6eb13f65bea894dbe7739b2e883d40a3e ]

    Downstream Port Containment (PCIe r5.0, sec. 6.2.10) disables the link upon
    an error and attempts to re-enable it when instructed by the DPC driver.

    A slot which is both DPC- and hotplug-capable is currently powered off by
    pciehp once DPC is triggered (due to the link change) and powered back up
    on successful recovery.  That's undesirable, the slot should remain powered
    so the hotplugged device remains bound to its driver.  DPC notifies the
    driver of the error and of successful recovery in pcie_do_recovery() and
    the driver may then restore the device to working state.

    Moreover, Sinan points out that turning off slot power by pciehp may foil
    recovery by DPC:  Power off/on is a cold reset concurrently to DPC's warm
    reset.  Sathyanarayanan reports extended delays or failure in link
    retraining by DPC if pciehp brings down the slot.

    Fix by detecting whether a Link Down event is caused by DPC and awaiting
    recovery if so.  On successful recovery, ignore both the Link Down and the
    subsequent Link Up event.

    Afterwards, check whether the link is down to detect surprise-removal or
    another DPC event immediately after DPC recovery.  Ensure that the
    corresponding DLLSC event is not ignored by synthesizing it and invoking
    irq_wake_thread() to trigger a re-run of pciehp_ist().

    The IRQ threads of the hotplug and DPC drivers, pciehp_ist() and
    dpc_handler(), race against each other.  If pciehp is faster than DPC, it
    will wait until DPC recovery completes.

    Recovery consists of two steps:  The first step (waiting for link
    disablement) is recognizable by pciehp through a set DPC Trigger Status
    bit.  The second step (waiting for link retraining) is recognizable through
    a newly introduced PCI_DPC_RECOVERING flag.
```

```
    If DPC is faster than pciehp, neither of the two flags will be set and
    pciehp may glean the recovery status from the new PCI_DPC_RECOVERED flag.
    The flag is zero if DPC didn't occur at all, hence DLLSC events are not
    ignored by default.

    pciehp waits up to 4 seconds before assuming that DPC recovery failed and
    bringing down the slot.  This timeout is not taken from the spec (it
    doesn't mandate one) but based on a report from Yicong Yang that DPC may
    take a bit more than 3 seconds on HiSilicon's Kunpeng platform.

    The timeout is necessary because the DPC Trigger Status bit may never
    clear:  On Root Ports which support RP Extensions for DPC, the DPC driver
    polls the DPC RP Busy bit for up to 1 second before giving up on DPC
    recovery.  Without the timeout, pciehp would then wait indefinitely for DPC
    to complete.

    This commit draws inspiration from previous attempts to synchronize DPC
    with pciehp:

    By Sinan Kaya, August 2018:
    https://lore.kernel.org/linux-pci/20180818065126.77912-1-okaya@kernel.org/

    By Ethan Zhao, October 2020:
    https://lore.kernel.org/linux-pci/20201007113158.48933-1-haifeng.zhao@intel.com/

    By Kuppuswamy Sathyanarayanan, March 2021:
    https://lore.kernel.org/linux-
pci/59cb30f5e5ac6d65427ceaadf1012b2ba8dbf66c.1615606143.git.sathyanarayanan.kuppuswamy@linux
.intel.com/

    Link:
https://lore.kernel.org/r/0be565d97438fe2a6d57354b3aa4e8626952a00b.1619857124.git.lukas@wunn
er.de
    Reported-by: Sinan Kaya <okaya@kernel.org>
    Reported-by: Ethan Zhao <haifeng.zhao@intel.com>
    Reported-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Tested-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Tested-by: Yicong Yang <yangyicong@hisilicon.com>
    Signed-off-by: Lukas Wunner <lukas@wunner.de>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Cc: Dan Williams <dan.j.williams@intel.com>
    Cc: Ashok Raj <ashok.raj@intel.com>
    Cc: Keith Busch <kbusch@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9f397e9e497d94e054db9a2cfdd6503fc65ad1b2
Author: Trond Myklebust <trond.myklebust@hammerspace.com>
Date:   Sat May 8 10:01:32 2021 -0400

    NFSv4: Fix delegation return in cases where we have to retry

    [ Upstream commit be20037725d17935ec669044bd2b15bc40c3b5ab ]

    If we're unable to immediately recover all locks because the server is
    unable to immediately service our reclaim calls, then we want to retry
    after we've finished servicing all the other asynchronous delegation
    returns on our queue.

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

commit eeb8022576c184b79a6a258ff77b78f263871b11
Author: Logan Gunthorpe <logang@deltatee.com>
Date:   Thu Jun 10 10:06:09 2021 -0600
```

```
    PCI/P2PDMA: Avoid pci_get_slot(), which may sleep

    [ Upstream commit 3ec0c3ec2d92c09465534a1ff9c6f9d9506ffef6 ]

    In order to use upstream_bridge_distance_warn() from a dma_map function, it
    must not sleep. However, pci_get_slot() takes the pci_bus_sem so it might
    sleep.

    In order to avoid this, try to get the host bridge's device from the first
    element in the device list. It should be impossible for the host bridge's
    device to go away while references are held on child devices, so the first
    element should not be able to change and, thus, this should be safe.

    Introduce a static function called pci_host_bridge_dev() to obtain the host
    bridge's root device.

    Link: https://lore.kernel.org/r/20210610160609.28447-7-logang@deltatee.com
    Signed-off-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00fe7b64bba873326f4d8681108490d589056b23
Author: Nick Desaulniers <ndesaulniers@google.com>
Date:   Tue Jun 1 20:29:26 2021 +0100

    ARM: 9087/1: kprobes: test-thumb: fix for LLVM_IAS=1

    [ Upstream commit 8b95a7d90ce8160ac5cffd5bace6e2eba01a871e ]

    There's a few instructions that GAS infers operands but Clang doesn't;
    from what I can tell the Arm ARM doesn't say these are optional.

    F5.1.257 TBB, TBH T1 Halfword variant
    F5.1.238 STREXD T1 variant
    F5.1.84 LDREXD T1 variant

    Link: https://github.com/ClangBuiltLinux/linux/issues/1309

    Signed-off-by: Nick Desaulniers <ndesaulniers@google.com>
    Reviewed-by: Jian Cai <jiancai@google.com>
    Signed-off-by: Russell King <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a67dc65f59ed62bce3b3a5f8fa13b4daa9592d6f
Author: Bixuan Cui <cuibixuan@huawei.com>
Date:   Sat May 8 11:14:59 2021 +0800

    power: reset: gpio-poweroff: add missing MODULE_DEVICE_TABLE

    [ Upstream commit ed3443fb4df4e140a22f65144546c8a8e1e27f4e ]

    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>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00a34a993ac382dd669158a2d500e7340253679a
Author: Krzysztof Kozlowski <krzk@kernel.org>
Date:   Wed May 26 13:20:36 2021 -0400
```

```
    power: supply: max17040: Do not enforce (incorrect) interrupt trigger type

    [ Upstream commit 8bb2314fc22628333d89df83d695ff9a8d2a6eac ]

    Interrupt line can be configured on different hardware in different way,
    even inverted.  Therefore driver should not enforce specific trigger
    type - edge falling - but instead rely on Devicetree to configure it.

    The Maxim 14577/77836 datasheets describe the interrupt line as active
    low with a requirement of acknowledge from the CPU therefore the edge
    falling is not correct.

    Signed-off-by: Krzysztof Kozlowski <krzk@kernel.org>
    Acked-by: Iskren Chernev <iskren.chernev@gmail.com>
    Acked-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d40b03031a0c9b6a2d551f149e8d51a4d65fe83
Author: Krzysztof Kozlowski <krzk@kernel.org>
Date:   Wed May 26 13:20:35 2021 -0400

    power: supply: max17042: Do not enforce (incorrect) interrupt trigger type

    [ Upstream commit 7fbf6b731bca347700e460d94b130f9d734b33e9 ]

    Interrupt line can be configured on different hardware in different way,
    even inverted.  Therefore driver should not enforce specific trigger
    type - edge falling - but instead rely on Devicetree to configure it.

    The Maxim 17047/77693 datasheets describe the interrupt line as active
    low with a requirement of acknowledge from the CPU therefore the edge
    falling is not correct.

    The interrupt line is shared between PMIC and RTC driver, so using level
    sensitive interrupt is here especially important to avoid races.  With
    an edge configuration in case if first PMIC signals interrupt followed
    shortly after by the RTC, the interrupt might not be yet cleared/acked
    thus the second one would not be noticed.

    Signed-off-by: Krzysztof Kozlowski <krzk@kernel.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 471c250065e15b224f1e728fdb96edc47437ac57
Author: Clemens Gruber <clemens.gruber@pqgruber.com>
Date:   Fri May 7 15:18:44 2021 +0200

    pwm: pca9685: Restrict period change for enabled PWMs

    [ Upstream commit 6d6e7050276d40b5de97aa950d5d71057f2e2a25 ]

    Previously, the last used PWM channel could change the global prescale
    setting, even if other channels are already in use.

    Fix it by only allowing the first enabled PWM to change the global
    chip-wide prescale setting. If there is more than one channel in use,
    the prescale settings resulting from the chosen periods must match.

    GPIOs do not count as enabled PWMs as they are not using the prescaler
    and can't change it.

    Signed-off-by: Clemens Gruber <clemens.gruber@pqgruber.com>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit b050025f8de42a5c7cca5a81107d18bdfe23e870
Author: Long Li <longli@microsoft.com>
Date:   Wed May 12 01:06:40 2021 -0700

    PCI: hv: Fix a race condition when removing the device

    [ Upstream commit 94d22763207ac6633612b8d8e0ca4fba0f7aa139 ]

    On removing the device, any work item (hv_pci_devices_present() or
    hv_pci_eject_device()) scheduled on workqueue hbus->wq may still be running
    and race with hv_pci_remove().

    This can happen because the host may send PCI_EJECT or PCI_BUS_RELATIONS(2)
    and decide to rescind the channel immediately after that.

    Fix this by flushing/destroying the workqueue of hbus before doing hbus remove.

    Link: https://lore.kernel.org/r/1620806800-30983-1-git-send-email-
longli@linuxonhyperv.com
    Signed-off-by: Long Li <longli@microsoft.com>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Michael Kelley <mikelley@microsoft.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6b8813b09fababd3fb602bda0f5dd7f458f1c3d6
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun May 23 00:50:42 2021 +0200

    power: supply: ab8500: Enable USB and AC

    [ Upstream commit f9184a228d7a60ad56b810d549a7debb355f1be6 ]

    The vendor code tree supplies platform data to enable he
    USB charging for AB8500 and AB8500 and disable the AC
    charging on the AB8505. This was missed when the driver
    was submitted to the mainline kernel.

    Fix this by doing what the vendor kernel does: always
    register the USB charger, do not register the AC charger
    on the AB8505.

    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa092b2f5bb3121badbbbdb59c0791189eb54db6
Author: Linus Walleij <linus.walleij@linaro.org>
Date:   Sun May 23 00:50:41 2021 +0200

    power: supply: ab8500: Avoid NULL pointers

    [ Upstream commit 5bcb5087c9dd3dca1ff0ebd8002c5313c9332b56 ]

    Sometimes the code will crash because we haven't enabled
    AC or USB charging and thus not created the corresponding
    psy device. Fix it by checking that it is there before
    notifying.

    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 341db343768bc44f3512facc464021730d64071c
Author: Linus Walleij <linus.walleij@linaro.org>
```

```
Date:   Sun May 23 00:50:39 2021 +0200

    power: supply: ab8500: Move to componentized binding

    [ Upstream commit 1c1f13a006ed0d71bb5664c8b7e3e77a28da3beb ]

    The driver has problems with the different components of
    the charging code racing with each other to probe().

    This results in all four subdrivers populating battery
    information to ascertain that it is populated for their
    own needs for example.

    Fix this by using component probing and thus expressing
    to the kernel that these are dependent components.
    The probes can happen in any order and will only acquire
    resources such as state container, regulators and
    interrupts and initialize the data structures, but no
    execution happens until the .bind() callback is called.

    The charging driver is the main component and binds
    first, then bind in order the three subcomponents:
    ab8500-fg, ab8500-btemp and ab8500-chargalg.

    Do some housekeeping while we are moving the code around.
    Like use devm_* for IRQs so as to cut down on some
    boilerplate.

    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dc72a15859b2e604abb8a4bff123fbac8a0be92a
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Mon May 17 16:41:17 2021 -0700

    PCI: ftpci100: Rename macro name collision

    [ Upstream commit 5be967d5016ac5ffb9c4d0df51b48441ee4d5ed1 ]

    PCI_IOSIZE is defined in mach-loongson64/spaces.h, so change the name
    of the PCI_* macros in pci-ftpci100.c to use FTPCI_* so that they are
    more localized and won't conflict with other drivers or arches.

    ../drivers/pci/controller/pci-ftpci100.c:37: warning: "PCI_IOSIZE" redefined
       37 | #define PCI_IOSIZE 0x00
          |
    In file included from ../arch/mips/include/asm/addrspace.h:13,
    ...              from ../drivers/pci/controller/pci-ftpci100.c:15:
    arch/mips/include/asm/mach-loongson64/spaces.h:11: note: this is the location of the
previous definition
       11 | #define PCI_IOSIZE SZ_16M

    Suggested-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20210517234117.3660-1-rdunlap@infradead.org
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Lorenzo Pieralisi <lorenzo.pieralisi@arm.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Cc: Jiaxun Yang <jiaxun.yang@flygoat.com>
    Cc: Linus Walleij <linus.walleij@linaro.org>
    Cc: Krzysztof Wilczyński <kw@linux.com>
    Cc: Thomas Bogendoerfer <tsbogend@alpha.franken.de>
    Cc: linux-mips@vger.kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 6606e5a6811a30491195cf1495567c34fab18ae0
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Wed Apr 28 11:05:24 2021 +0200

    pwm: spear: Don't modify HW state in .remove callback

    [ Upstream commit b601a18f12383001e7a8da238de7ca1559ebc450 ]

    A consumer is expected to disable a PWM before calling pwm_put(). And if
    they didn't there is hopefully a good reason (or the consumer needs
    fixing). Also if disabling an enabled PWM was the right thing to do,
    this should better be done in the framework instead of in each low level
    driver.

    So drop the hardware modification from the .remove() callback.

    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Signed-off-by: Thierry Reding <thierry.reding@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19ce06dcd189279029fd2a3537c43f078bf96451
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 14:17:12 2021 +0800

    power: supply: sc2731_charger: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 2aac79d14d76879c8e307820b31876e315b1b242 ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d7dae32b14bf3c48599816470d0e0b2915f8327
Author: Zou Wei <zou_wei@huawei.com>
Date:   Tue May 11 11:37:45 2021 +0800

    power: supply: sc27xx: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 603fcfb9d4ec1cad8d66d3bb37f3613afa8a661a ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 77637270800c7df8d439127854d99506796cd5fa
Author: Marco Elver <elver@google.com>
Date:   Wed Jun 30 18:56:49 2021 -0700

    kcov: add __no_sanitize_coverage to fix noinstr for all architectures

    [ Upstream commit 540540d06e9d9b3769b46d88def90f7e7c002322 ]

    Until now no compiler supported an attribute to disable coverage
    instrumentation as used by KCOV.
```

```
    To work around this limitation on x86, noinstr functions have their
    coverage instrumentation turned into nops by objtool.  However, this
    solution doesn't scale automatically to other architectures, such as
    arm64, which are migrating to use the generic entry code.

    Clang [1] and GCC [2] have added support for the attribute recently.
    [1] https://github.com/llvm/llvm-project/commit/280333021e9550d80f5c1152a34e33e81df1e178
    [2] https://gcc.gnu.org/git/?
p=gcc.git;a=commit;h=cec4d4a6782c9bd8d071839c50a239c49caca689
    The changes will appear in Clang 13 and GCC 12.

    Add __no_sanitize_coverage for both compilers, and add it to noinstr.

    Note: In the Clang case, __has_feature(coverage_sanitizer) is only true if
    the feature is enabled, and therefore we do not require an additional
    defined(CONFIG_KCOV) (like in the GCC case where __has_attribute(..) is
    always true) to avoid adding redundant attributes to functions if KCOV is
    off.  That being said, compilers that support the attribute will not
    generate errors/warnings if the attribute is redundantly used; however,
    where possible let's avoid it as it reduces preprocessed code size and
    associated compile-time overheads.

    [elver@google.com: Implement __has_feature(coverage_sanitizer) in Clang]
      Link: https://lkml.kernel.org/r/20210527162655.3246381-1-elver@google.com
    [elver@google.com: add comment explaining __has_feature() in Clang]
      Link: https://lkml.kernel.org/r/20210527194448.3470080-1-elver@google.com

    Link: https://lkml.kernel.org/r/20210525175819.699786-1-elver@google.com
    Signed-off-by: Marco Elver <elver@google.com>
    Acked-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Miguel Ojeda <ojeda@kernel.org>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Will Deacon <will@kernel.org>
    Cc: Ard Biesheuvel <ardb@kernel.org>
    Cc: Luc Van Oostenryck <luc.vanoostenryck@gmail.com>
    Cc: Arvind Sankar <nivedita@alum.mit.edu>
    Cc: Masahiro Yamada <masahiroy@kernel.org>
    Cc: Sami Tolvanen <samitolvanen@google.com>
    Cc: Arnd Bergmann <arnd@arndb.de>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: Mark Rutland <mark.rutland@arm.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 67a8dfc86bdb21c8f01bb4a5e5f9c89a063807b0
Author: Dimitri John Ledkov <dimitri.ledkov@canonical.com>
Date:   Wed Jun 30 18:56:16 2021 -0700

    lib/decompress_unlz4.c: correctly handle zero-padding around initrds.

    [ Upstream commit 2c484419efc09e7234c667aa72698cb79ba8d8ed ]

    lz4 compatible decompressor is simple.  The format is underspecified and
    relies on EOF notification to determine when to stop.  Initramfs buffer
    format[1] explicitly states that it can have arbitrary number of zero
    padding.  Thus when operating without a fill function, be extra careful to
    ensure that sizes less than 4, or apperantly empty chunksizes are treated
    as EOF.

    To test this I have created two cpio initrds, first a normal one,
    main.cpio.  And second one with just a single /test-file with content
```

```
    "second" second.cpio.  Then i compressed both of them with gzip, and with
    lz4 -l.  Then I created a padding of 4 bytes (dd if=/dev/zero of=pad4 bs=1
    count=4).  To create four testcase initrds:

     1) main.cpio.gzip + extra.cpio.gzip = pad0.gzip
     2) main.cpio.lz4  + extra.cpio.lz4 = pad0.lz4
     3) main.cpio.gzip + pad4 + extra.cpio.gzip = pad4.gzip
     4) main.cpio.lz4  + pad4 + extra.cpio.lz4 = pad4.lz4

    The pad4 test-cases replicate the initrd load by grub, as it pads and
    aligns every initrd it loads.

    All of the above boot, however /test-file was not accessible in the initrd
    for the testcase #4, as decoding in lz4 decompressor failed.  Also an
    error message printed which usually is harmless.

    Whith a patched kernel, all of the above testcases now pass, and
    /test-file is accessible.

    This fixes lz4 initrd decompress warning on every boot with grub.  And
    more importantly this fixes inability to load multiple lz4 compressed
    initrds with grub.  This patch has been shipping in Ubuntu kernels since
    January 2021.

    [1] ./Documentation/driver-api/early-userspace/buffer-format.rst

    BugLink: https://bugs.launchpad.net/bugs/1835660
    Link: https://lore.kernel.org/lkml/20210114200256.196589-1-xnox@ubuntu.com/ # v0
    Link: https://lkml.kernel.org/r/20210513104831.432975-1-dimitri.ledkov@canonical.com
    Signed-off-by: Dimitri John Ledkov <dimitri.ledkov@canonical.com>
    Cc: Kyungsik Lee <kyungsik.lee@lge.com>
    Cc: Yinghai Lu <yinghai@kernel.org>
    Cc: Bongkyu Kim <bongkyu.kim@lge.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Sven Schmidt <4sschmid@informatik.uni-hamburg.de>
    Cc: Rajat Asthana <thisisrast7@gmail.com>
    Cc: Nick Terrell <terrelln@fb.com>
    Cc: Gao Xiang <hsiangkao@redhat.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit af6365ae0991c304387cc0297beb55974f67a29f
Author: Rashmi A <rashmi.a@intel.com>
Date:   Thu Jun 3 23:52:42 2021 +0530

    phy: intel: Fix for warnings due to EMMC clock 175Mhz change in FIP

    [ Upstream commit 2f2b73a29d2aabf5ad0150856c3e5cb6e04dcfc1 ]

    Since the EMMC clock was changed from 200Mhz to 175Mhz in FIP,
    there were some warnings introduced, as the frequency values
    being checked was still wrt 200Mhz in code. Hence, the frequency
    checks are now updated based on the current 175Mhz EMMC clock changed
    in FIP.

    Spamming kernel log msg:
    "phy phy-20290000.mmc_phy.2: Unsupported rate: 43750000"

    Signed-off-by: Rashmi A <rashmi.a@intel.com>
    Reviewed-by: Adrian Hunter <adrian.hunter@intel.com>
    Acked-By: Vinod Koul <vkoul@kernel.org>
    Link: https://lore.kernel.org/r/20210603182242.25733-3-rashmi.a@intel.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit e00d3efd53d0f9f63cd73b4bbf8888da944e1983
Author: Dmitry Torokhov <dmitry.torokhov@gmail.com>
Date:   Fri Jun 4 16:27:44 2021 -0700

    i2c: core: Disable client irq on reboot/shutdown

    [ Upstream commit b64210f2f7c11c757432ba3701d88241b2b98fb1 ]

    If an i2c client receives an interrupt during reboot or shutdown it may
    be too late to service it by making an i2c transaction on the bus
    because the i2c controller has already been shutdown. This can lead to
    system hangs if the i2c controller tries to make a transfer that is
    doomed to fail because the access to the i2c pins is already shut down,
    or an iommu translation has been torn down so i2c controller register
    access doesn't work.

    Let's simply disable the irq if there isn't a shutdown callback for an
    i2c client when there is an irq associated with the device. This will
    make sure that irqs don't come in later than the time that we can handle
    it. We don't do this if the i2c client device already has a shutdown
    callback because presumably they're doing the right thing and quieting
    the device so irqs don't come in after the shutdown callback returns.

    Reported-by: kernel test robot <lkp@intel.com>
    [swboyd@chromium.org: Dropped newline, added commit text, added
    interrupt.h for robot build error]
    Signed-off-by: Stephen Boyd <swboyd@chromium.org>
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Wolfram Sang <wsa@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47419e67a5ff6d5c97956f1135f448bce89f3c6e
Author: Alexander Shishkin <alexander.shishkin@linux.intel.com>
Date:   Mon Jun 21 18:12:46 2021 +0300

    intel_th: Wait until port is in reset before programming it

    [ Upstream commit ab1afed701d2db7eb35c1a2526a29067a38e93d1 ]

    Some devices don't drain their pipelines if we don't make sure that
    the corresponding output port is in reset before programming it for
    a new trace capture, resulting in bits of old trace appearing in the
    new trace capture. Fix that by explicitly making sure the reset is
    asserted before programming new trace capture.

    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Alexander Shishkin <alexander.shishkin@linux.intel.com>
    Link: https://lore.kernel.org/r/20210621151246.31891-5-
alexander.shishkin@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8c012f3504abad245ff753fb89d42427adb9e7ff
Author: Fabio Aiuto <fabioaiuto83@gmail.com>
Date:   Tue Jun 22 15:10:00 2021 +0200

    staging: rtl8723bs: fix check allowing 5Ghz settings

    [ Upstream commit 990a1472930bf2bb7927ea2def4b434790780a8d ]

    fix check allowing 5Ghz settings, only disabled and
    2.4Ghz enabled states are allowed. Fix comment
    accordingly.
```

```
    Acked-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Fabio Aiuto <fabioaiuto83@gmail.com>
    Link:
https://lore.kernel.org/r/df7d0ecc02ac7a27e568768523dd7b3f34acd551.1624367072.git.fabioaiuto
83@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 32615857345b7bab16a79b905e8eef1001bd4a4b
Author: Fabio Aiuto <fabioaiuto83@gmail.com>
Date:   Tue Jun 22 15:09:58 2021 +0200

    staging: rtl8723bs: fix macro value for 2.4Ghz only device

    [ Upstream commit 6d490a27e23c5fb79b766530016ab8665169498e ]

    fix IQK_Matrix_Settings_NUM macro value to 14 which is
    the max channel number value allowed in a 2.4Ghz device.

    Acked-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Fabio Aiuto <fabioaiuto83@gmail.com>
    Link:
https://lore.kernel.org/r/0b4a876929949248aa18cb919da3583c65e4ee4e.1624367072.git.fabioaiuto
83@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f59df2b2358122e275b7e27af74fddd68c7e3f7
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 14:49:18 2021 +0800

    leds: turris-omnia: add missing MODULE_DEVICE_TABLE

    [ Upstream commit 9d0150db97583cfbb6b44cbe02241a1a48f90210 ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Pavel Machek <pavel@ucw.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 658d9b2911933bfaf78b53d08313b6dc88e8670a
Author: Takashi Sakamoto <o-takashi@sakamocchi.jp>
Date:   Wed Jun 23 16:59:33 2021 +0900

    ALSA: firewire-motu: fix detection for S/PDIF source on optical interface in v2 protocol

    [ Upstream commit fa4db23233eb912234bdfb0b26a38be079c6b5ea ]

    The devices in protocol version 2 has a register with flag for IEC 60958
    signal detection as source of sampling clock without discrimination
    between coaxial and optical interfaces. On the other hand, current
    implementation of driver manage to interpret type of signal on optical
    interface instead.

    This commit fixes the detection of optical/coaxial interface for S/PDIF
    signal.

    Signed-off-by: Takashi Sakamoto <o-takashi@sakamocchi.jp>
    Link: https://lore.kernel.org/r/20210623075941.72562-2-o-takashi@sakamocchi.jp
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 048d83c67d37378b2200afa9a494bf75cec07307
Author: Geoffrey D. Bennett <g@b4.vu>
Date:   Tue Jun 22 03:39:18 2021 +0930

    ALSA: usb-audio: scarlett2: Fix 6i6 Gen 2 line out descriptions

    [ Upstream commit c712c6c0ff2d60478582e337185bcdd520a7dc2e ]

    There are two headphone outputs, and they map to the four analogue
    outputs.

    Signed-off-by: Geoffrey D. Bennett <g@b4.vu>
    Link:
https://lore.kernel.org/r/205e5e5348f08ded0cc4da5446f604d4b91db5bf.1624294591.git.g@b4.vu
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4348f36461b840697ee9652109b24921afa3d38b
Author: Jiajun Cao <jjcao20@fudan.edu.cn>
Date:   Tue Jun 22 21:19:42 2021 +0800

    ALSA: hda: Add IRQ check for platform_get_irq()

    [ Upstream commit 8c13212443230d03ff25014514ec0d53498c0912 ]

    The function hda_tegra_first_init() neglects to check the return
    value after executing platform_get_irq().

    hda_tegra_first_init() should check the return value (if negative
    error number) for errors so as to not pass a negative value to
    the devm_request_irq().

    Fix it by adding a check for the return value irq_id.

    Signed-off-by: Jiajun Cao <jjcao20@fudan.edu.cn>
    Signed-off-by: Xin Tan <tanxin.ctf@gmail.com>
    Reviewed-by: Thierry Reding <treding@nvidia.com>
    Link: https://lore.kernel.org/r/20210622131947.94346-1-jjcao20@fudan.edu.cn
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eff67a0bb38ffb3ad94fdc7f4a930eed5c9e173b
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Mon Jun 21 14:21:47 2021 +0200

    backlight: lm3630a: Fix return code of .update_status() callback

    [ Upstream commit b9481a667a90ec739995e85f91f3672ca44d6ffa ]

    According to <linux/backlight.h> .update_status() is supposed to
    return 0 on success and a negative error code otherwise. Adapt
    lm3630a_bank_a_update_status() and lm3630a_bank_b_update_status() to
    actually do it.

    While touching that also add the error code to the failure message.

    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 18fbea65d1a52dc64283cb3122924e6d9fae00d0
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Jun 21 14:40:52 2021 -0500
```

```
    ASoC: Intel: kbl_da7219_max98357a: shrink platform_id below 20 characters

    [ Upstream commit 94efd726b947f265bd313605c9f73edec5469d65 ]

    Sparse throws the following warnings:

    sound/soc/intel/boards/kbl_da7219_max98357a.c:647:25: error: too long
    initializer-string for array of char(no space for nul char)

    Fix by using the 'mx' acronym for Maxim.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Paul Olaru <paul.olaru@oss.nxp.com>
    Reviewed-by: Guennadi Liakhovetski <guennadi.liakhovetski@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Link: https://lore.kernel.org/r/20210621194057.21711-6-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d422f8be38a224fdaeec97f930a5c6f098a6ffba
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Jun 15 09:39:22 2021 +0800

    ASoC: fsl_xcvr: check return value after calling platform_get_resource_byname()

    [ Upstream commit a2f6ed4a44721d3a9fdf4da7e0743cb13866bf61 ]

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

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20210615013922.784296-10-yangyingliang@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d15c57f8f120bef8343e01de84ddb8d094bb1f1f
Author: Benjamin Herrenschmidt <benh@kernel.crashing.org>
Date:   Fri Jun 18 13:49:00 2021 +1000

    powerpc/boot: Fixup device-tree on little endian

    [ Upstream commit c93f80849bdd9b45d834053ae1336e28f0026c84 ]

    This fixes the core devtree.c functions and the ns16550 UART backend.

    Signed-off-by: Benjamin Herrenschmidt <benh@kernel.crashing.org>
    Signed-off-by: Paul Mackerras <paulus@ozlabs.org>
    Reviewed-by: Segher Boessenkool <segher@kernel.crashing.org>
    Acked-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/YMwXrPT8nc4YUdJ9@thinks.paulus.ozlabs.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 34ada7b357dda19b32a1c440512da5367621be20
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Fri Jun 18 12:38:35 2021 +0800

    usb: gadget: hid: fix error return code in hid_bind()

    [ Upstream commit 88693f770bb09c196b1eb5f06a484a254ecb9924 ]

    Fix to return a negative error code from the error handling
    case instead of 0.
```

```
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20210618043835.2641360-1-yangyingliang@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c904d51536b436f709de7d188ec15d8ffd02c1ff
Author: Ruslan Bilovol <ruslan.bilovol@gmail.com>
Date:   Thu Jun 17 19:27:55 2021 +0300

    usb: gadget: f_hid: fix endianness issue with descriptors

    [ Upstream commit 33cb46c4676d01956811b68a29157ea969a5df70 ]

    Running sparse checker it shows warning message about
    incorrect endianness used for descriptor initialization:

    | f_hid.c:91:43: warning: incorrect type in initializer (different base types)
    | f_hid.c:91:43:    expected restricted __le16 [usertype] bcdHID
    | f_hid.c:91:43:    got int

    Fixing issue with cpu_to_le16() macro, however this is not a real issue
    as the value is the same both endians.

    Cc: Fabien Chouteau <fabien.chouteau@barco.com>
    Cc: Segiy Stetsyuk <serg_stetsuk@ukr.net>
    Signed-off-by: Ruslan Bilovol <ruslan.bilovol@gmail.com>
    Link: https://lore.kernel.org/r/20210617162755.29676-1-ruslan.bilovol@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7113c2692f7fc0f237baf7b2b337cb27addc9c6a
Author: Geoffrey D. Bennett <g@b4.vu>
Date:   Mon Jun 21 02:16:45 2021 +0930

    ALSA: usb-audio: scarlett2: Fix scarlett2_*_ctl_put() return values

    [ Upstream commit c5d8e008032f3cd5f266d552732973a960b0bd4b ]

    Mixer control put callbacks should return 1 if the value is changed.
    Fix the sw_hw, level, pad, and button controls accordingly.

    Signed-off-by: Geoffrey D. Bennett <g@b4.vu>
    Link: https://lore.kernel.org/r/20210620164645.GA9221@m.b4.vu
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d94bcba9e59843622106027a489852db63f49b8
Author: Geoffrey D. Bennett <g@b4.vu>
Date:   Mon Jun 21 02:16:43 2021 +0930

    ALSA: usb-audio: scarlett2: Fix data_mutex lock

    [ Upstream commit 9b5ddea9ce5a68d7d2bedcb69901ac2a86c96c7b ]

    The private->vol_updated flag was being checked outside of the
    mutex_lock/unlock() of private->data_mutex leading to the volume data
    being fetched twice from the device unnecessarily or old volume data
    being returned.

    Update scarlett2_*_ctl_get() and include the private->vol_updated flag
    check inside the critical region.

    Signed-off-by: Geoffrey D. Bennett <g@b4.vu>
```

```
    Link: https://lore.kernel.org/r/20210620164643.GA9216@m.b4.vu
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f3f7f6ef23fc2f1df2efeda5c2ea4e5d7423546a
Author: Geoffrey D. Bennett <g@b4.vu>
Date:   Mon Jun 21 02:16:25 2021 +0930

    ALSA: usb-audio: scarlett2: Fix 18i8 Gen 2 PCM Input count

    [ Upstream commit c5210f213456383482b4a77c5310282a89a106b5 ]

    The 18i8 Gen 2 has 8 PCM Inputs, not 20. Fix the ports entry in
    s18i8_gen2_info.

    Signed-off-by: Geoffrey D. Bennett <g@b4.vu>
    Link: https://lore.kernel.org/r/20210620164625.GA9165@m.b4.vu
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 669eb773678163aa1f7bc9a645f6b88ac284b0b0
Author: Greg Ungerer <gerg@linux-m68k.org>
Date:   Tue Apr 27 23:49:31 2021 +1000

    m68knommu: fix missing LCD splash screen data initializer

    [ Upstream commit 3f605ee17f8e785ba6ff69bee8e584e377a3cf28 ]

    The kernel test robot reports that the "screen_bits" splash screen data
    is missing for the dragen platform:

        arch/m68k/68000/dragen2.c:73:16: error: 'screen_bits' undeclared (first use in this
function)
          73 |  LSSA = (long) screen_bits;
             |                ^~~~~~~~~~~
        arch/m68k/68000/dragen2.c:73:16: note: each undeclared identifier is reported only
once for each function it appears in

    Digging around a bit I found the screen_bits data structure was originally
    in a screen.h file that was generated from a screen.xbm file. That was
    removed in commit 0c0e6db80683 ("m68k: drop unused parts of 68VZ328
    Makefile").

    Other splash screen initializers for 68000 targets are kept here as the
    C data structures so lets do the same for this one. Add the missing
    screen.h file and include it in the dragen specific startup code.

    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Greg Ungerer <gerg@linux-m68k.org>
    Acked-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit efd33d33dedd9df9d711ca0535651d2741a645dd
Author: Takashi Sakamoto <o-takashi@sakamocchi.jp>
Date:   Sat Jun 19 17:39:22 2021 +0900

    ALSA: bebob: add support for ToneWeal FW66

    [ Upstream commit 50ebe56222bfa0911a932930f9229ee5995508d9 ]

    A user of FFADO project reported the issue of ToneWeal FW66. As a result,
    the device is identified as one of applications of BeBoB solution.

    I note that in the report the device returns contradictory result in plug
    discovery process for audio subunit. Fortunately ALSA BeBoB driver doesn't
```

```
    perform it thus it's likely to handle the device without issues.

    I receive no reaction to test request for this patch yet, however it would
    be worth to add support for it.

    daniel@gibbonmoon:/sys/bus/firewire/devices/fw1$ grep -r . *
    Binary file config_rom matches
    dev:244:1
    guid:0x0023270002000000
    hardware_version:0x000002
    is_local:0
    model:0x020002
    model_name:FW66
    power/runtime_active_time:0
    power/runtime_active_kids:0
    power/runtime_usage:0
    power/runtime_status:unsupported
    power/async:disabled
    power/runtime_suspended_time:0
    power/runtime_enabled:disabled
    power/control:auto
    subsystem/drivers_autoprobe:1
    uevent:MAJOR=244
    uevent:MINOR=1
    uevent:DEVNAME=fw1
    units:0x00a02d:0x010001
    vendor:0x002327
    vendor_name:ToneWeal
    fw1.0/uevent:MODALIAS=ieee1394:ven00002327mo00020002sp0000A02Dver00010001
    fw1.0/power/runtime_active_time:0
    fw1.0/power/runtime_active_kids:0
    fw1.0/power/runtime_usage:0
    fw1.0/power/runtime_status:unsupported
    fw1.0/power/async:disabled
    fw1.0/power/runtime_suspended_time:0
    fw1.0/power/runtime_enabled:disabled
    fw1.0/power/control:auto
    fw1.0/model:0x020002
    fw1.0/rom_index:15
    fw1.0/specifier_id:0x00a02d
    fw1.0/model_name:FW66
    fw1.0/version:0x010001
    fw1.0/modalias:ieee1394:ven00002327mo00020002sp0000A02Dver00010001

    Cc: Daniel Jozsef <daniel.jozsef@gmail.com>
    Reference: https://lore.kernel.org/alsa-devel/20200119164335.GA11974@workstation/
    Signed-off-by: Takashi Sakamoto <o-takashi@sakamocchi.jp>
    Link: https://lore.kernel.org/r/20210619083922.16060-1-o-takashi@sakamocchi.jp
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5dc9f63b20d768df89edcdbe58966b1e4a1df0d
Author: Yizhuo Zhai <yzhai003@ucr.edu>
Date:   Sat Jun 19 22:26:50 2021 -0700

    Input: hideep - fix the uninitialized use in hideep_nvm_unlock()

    [ Upstream commit cac7100d4c51c04979dacdfe6c9a5e400d3f0a27 ]

    Inside function hideep_nvm_unlock(), variable "unmask_code" could
    be uninitialized if hideep_pgm_r_reg() returns error, however, it
    is used in the later if statement after an "and" operation, which
    is potentially unsafe.

    Signed-off-by: Yizhuo <yzhai003@ucr.edu>
```

```
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1043c57cff3d055866baa6a85dc06e593a582f8e
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Jun 14 22:09:20 2021 +0200

    s390/mem_detect: fix tprot() program check new psw handling

    [ Upstream commit da9057576785aaab52e706e76c0475c85b77ec14 ]

    The tprot() inline asm temporarily changes the program check new psw
    to redirect a potential program check on the diag instruction.
    Restoring of the program check new psw is done in C code behind the
    inline asm.

    This can be problematic, especially if the function is inlined, since
    the compiler can reorder instructions in such a way that a different
    instruction, which may result in a program check, might be executed
    before the program check new psw has been restored.

    To avoid such a scenario move restoring into the inline asm. For
    consistency reasons move also saving of the original program check new
    psw into the inline asm.

    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf04c373f1b30f0b443f6e36c33fa090cf3708c7
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Jun 14 21:40:43 2021 +0200

    s390/mem_detect: fix diag260() program check new psw handling

    [ Upstream commit 86807f348f418a84970eebb8f9912a7eea16b497 ]

    The __diag260() inline asm temporarily changes the program check new
    psw to redirect a potential program check on the diag instruction.
    Restoring of the program check new psw is done in C code behind the
    inline asm.

    This can be problematic, especially if the function is inlined, since
    the compiler can reorder instructions in such a way that a different
    instruction, which may result in a program check, might be executed
    before the program check new psw has been restored.

    To avoid such a scenario move restoring into the inline asm. For
    consistency reasons move also saving of the original program check new
    psw into the inline asm.

    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c80766b911d3d7f4326bba2fedf415064292cce
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Jun 14 14:49:31 2021 +0200

    s390/ipl_parm: fix program check new psw handling

    [ Upstream commit 88c2510cecb7e2b518e3c4fdf3cf0e13ebe9377c ]

    The __diag308() inline asm temporarily changes the program check new
    psw to redirect a potential program check on the diag instruction.
```

```
    Restoring of the program check new psw is done in C code behind the
    inline asm.

    This can be problematic, especially if the function is inlined, since
    the compiler can reorder instructions in such a way that a different
    instruction, which may result in a program check, might be executed
    before the program check new psw has been restored.

    To avoid such a scenario move restoring into the inline asm. For
    consistency reasons move also saving of the original program check new
    psw into the inline asm.

    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 45b658ee680d889ba2d5e825c064803851a9ad49
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Thu Jun 10 17:50:25 2021 +0200

    s390/processor: always inline stap() and __load_psw_mask()

    [ Upstream commit 9c9a915afd90f7534c16a71d1cd44b58596fddf3 ]

    s390 is the only architecture which makes use of the __no_kasan_or_inline
    attribute for two functions. Given that both stap() and __load_psw_mask()
    are very small functions they can and should be always inlined anyway.

    Therefore get rid of __no_kasan_or_inline and always inline these
    functions.

    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4de8aba5c10fde25e05826d2899fb187f6ed6dfa
Author: Koby Elbaz <kelbaz@habana.ai>
Date:   Thu Jun 10 09:14:43 2021 +0300

    habanalabs/gaudi: set the correct rc in case of err

    [ Upstream commit 1f7ef4bf41c7c2abad3d21b8c69db11fc3ebc4f5 ]

    fix the following smatch warnings:
    gaudi_internal_cb_pool_init() warn: missing error code 'rc'

    Signed-off-by: Koby Elbaz <kelbaz@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 29d01e4b4c04f6822f99b5b60d18c824aa011839
Author: Koby Elbaz <kelbaz@habana.ai>
Date:   Wed Jun 9 21:43:52 2021 +0300

    habanalabs: remove node from list before freeing the node

    [ Upstream commit f5eb7bf0c487a212ebda3c1b048fc3ccabacc147 ]

    fix the following smatch warnings:

    goya_pin_memory_before_cs()
    warn: '&userptr->job_node' not removed from list

    gaudi_pin_memory_before_cs()
```

```
      warn: '&userptr->job_node' not removed from list

      Signed-off-by: Koby Elbaz <kelbaz@habana.ai>
      Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit cce27a77cc59b8a9758b7c084927374566a6a01b
  Author: Koby Elbaz <kelbaz@habana.ai>
  Date:   Thu Jun 10 09:01:57 2021 +0300

      habanalabs: set rc as 'valid' in case of intentional func exit

      [ Upstream commit 11d5cb8b95456e2432dfee2ffcebf0623998493a ]

      fix the following smatch warnings:
      hl_fw_static_init_cpu() warn: missing error code 'rc'

      Signed-off-by: Koby Elbaz <kelbaz@habana.ai>
      Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 275674c1e58be13db5cf1a6439b5c024cbec8e0f
  Author: Ohad Sharabi <osharabi@habana.ai>
  Date:   Thu Jun 3 00:24:32 2021 +0300

      habanalabs: fix mask to obtain page offset

      [ Upstream commit 0f37510ca34848718db1003479bb4671e8f3c112 ]

      When converting virtual address to physical we need to add correct
      offset to the physical page.

      For this we need to use mask that include ALL bits of page offset.

      Signed-off-by: Ohad Sharabi <osharabi@habana.ai>
      Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 30dd0cbabf14a5e310cb11b2c547791a6431ad37
  Author: Koby Elbaz <kelbaz@habana.ai>
  Date:   Wed May 19 15:16:52 2021 +0300

      habanalabs/gaudi: set the correct cpu_id on MME2_QM failure

      [ Upstream commit b92c637c5f5ef7e3e21dbc7bfa7f1999450f3902 ]

      This fix was applied since there was an incorrect reported CPU ID to GIC
      such that an error in MME2 QMAN aliased to be an arriving from DMA0_QM.

      Signed-off-by: Koby Elbaz <kelbaz@habana.ai>
      Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4bf6fcdcca815f943562f8d500c9394829c8806b
  Author: Ohad Sharabi <osharabi@habana.ai>
  Date:   Tue May 11 16:02:41 2021 +0300

      habanalabs: check if asic secured with asic type

      [ Upstream commit 190ec49710a9fe0d5e9e36fe1a2fa864c048484f ]
```

```
    Fix issue in which the input to the function is_asic_secured was device
    PCI_IDS number instead of the asic_type enumeration.

    Signed-off-by: Ohad Sharabi <osharabi@habana.ai>
    Reviewed-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Oded Gabbay <ogabbay@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dac25ad9dd9d822f3dd705a30ad5655db83a3b6d
Author: Mathias Nyman <mathias.nyman@linux.intel.com>
Date:   Thu Jun 17 18:03:53 2021 +0300

    xhci: handle failed buffer copy to URB sg list and fix a W=1 copiler warning

    [ Upstream commit 271a21d8b280b186f8cc9ca6f7151902efde9512 ]

    Set the urb->actual_length to bytes successfully copied in case all bytes
    weren't copied from a temporary buffer to the URB sg list.
    Also print a debug message

    Signed-off-by: Mathias Nyman <mathias.nyman@linux.intel.com>
    Link: https://lore.kernel.org/r/20210617150354.1512157-4-mathias.nyman@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 44be4fd02a920c3b7c2cfb355bda64c45ace4c40
Author: Zhen Lei <thunder.leizhen@huawei.com>
Date:   Thu Jun 17 18:37:29 2021 +0800

    ASoC: soc-core: Fix the error return code in snd_soc_of_parse_audio_routing()

    [ Upstream commit 7d3865a10b9ff2669c531d5ddd60bf46b3d48f1e ]

    When devm_kcalloc() fails, the error code -ENOMEM should be returned
    instead of -EINVAL.

    Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
    Link: https://lore.kernel.org/r/20210617103729.1918-1-thunder.leizhen@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c90bcc03dd7c30603111a1ceb5f87d86c242d0b7
Author: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
Date:   Thu Jun 10 14:06:39 2021 +0530

    powerpc/mm/book3s64: Fix possible build error

    [ Upstream commit 07d8ad6fd8a3d47f50595ca4826f41dbf4f3a0c6 ]

    Update _tlbiel_pid() such that we can avoid build errors like below when
    using this function in other places.

    arch/powerpc/mm/book3s64/radix_tlb.c: In function ‘__radix__flush_tlb_range_psize’:
    arch/powerpc/mm/book3s64/radix_tlb.c:114:2: warning: ‘asm’ operand 3 probably does not
match constraints
      114 |  asm volatile(PPC_TLBIEL(%0, %4, %3, %2, %1)
          |  ^~~
    arch/powerpc/mm/book3s64/radix_tlb.c:114:2: error: impossible constraint in ‘asm’
    make[4]: *** [scripts/Makefile.build:271: arch/powerpc/mm/book3s64/radix_tlb.o] Error 1
    m

    With this fix, we can also drop the __always_inline in __radix_flush_tlb_range_psize
    which was added by commit e12d6d7d46a6 ("powerpc/mm/radix: mark
__radix__flush_tlb_range_psize() as __always_inline")
```

```
    Signed-off-by: Aneesh Kumar K.V <aneesh.kumar@linux.ibm.com>
    Reviewed-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Acked-by: Michael Ellerman <mpe@ellerman.id.au>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/20210610083639.387365-1-aneesh.kumar@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4921acaab931bb3f95440d58b55b2d3093aed34
Author: Peter Robinson <pbrobinson@gmail.com>
Date:   Sun Jun 13 23:03:26 2021 +0100

    gpio: pca953x: Add support for the On Semi pca9655

    [ Upstream commit 6d49b3a0f351925b5ea5047166c112b7590b918a ]

    The On Semi pca9655 is a 16 bit variant of the On Semi pca9654 GPIO
    expander, with 16 GPIOs and interrupt functionality.

    Signed-off-by: Peter Robinson <pbrobinson@gmail.com>
    [Bartosz: fixed indentation as noted by Andy]
    Signed-off-by: Bartosz Golaszewski <bgolaszewski@baylibre.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac2ee946b71ef521c5e612ec214f59b31bdf8f4e
Author: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
Date:   Tue May 25 09:51:42 2021 -0400

    selftests/powerpc: Fix "no_handler" EBB selftest

    [ Upstream commit 45677c9aebe926192e59475b35a1ff35ff2d4217 ]

    The "no_handler_test" in ebb selftests attempts to read the PMU
    registers twice via helper function "dump_ebb_state". First dump is
    just before closing of event and the second invocation is done after
    closing of the event. The original intention of second
    dump_ebb_state was to dump the state of registers at the end of
    the test when the counters are frozen. But this will be achieved
    with the first call itself since sample period is set to low value
    and PMU will be frozen by then. Hence patch removes the
    dump which was done before closing of the event.

    Reported-by: Shirisha Ganta <shirisha.ganta1@ibm.com>
    Signed-off-by: Athira Rajeev <atrajeev@linux.vnet.ibm.com>
    Tested-by: Nageswara R Sastry <rnsastry@linux.ibm.com <mailto:rnsastry@linux.ibm.com>>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://lore.kernel.org/r/1621950703-1532-2-git-send-email-
atrajeev@linux.vnet.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df6b21c0aa9d8bf88613f10638aca1e9a4b632fa
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Wed Jun 16 10:11:21 2021 +0800

    ALSA: ppc: fix error return code in snd_pmac_probe()

    [ Upstream commit 80b9c1be567c3c6bbe0d4b290af578e630485b5d ]

    If snd_pmac_tumbler_init() or snd_pmac_tumbler_post_init() fails,
    snd_pmac_probe() need return error code.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20210616021121.1991502-1-yangyingliang@huawei.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 21b28dfedfc118b88cfdce6d7e0d5bab362800b5
Author: Michael Kelley <mikelley@microsoft.com>
Date:   Fri Jun 4 10:21:03 2021 -0700

    scsi: storvsc: Correctly handle multiple flags in srb_status

    [ Upstream commit 52e1b3b3daa9d53f0204bf474ee1d4b1beb38234 ]

    Hyper-V is observed to sometimes set multiple flags in the srb_status, such
    as ABORTED and ERROR. Current code in storvsc_handle_error() handles only a
    single flag being set, and does nothing when multiple flags are set.  Fix
    this by changing the case statement into a series of "if" statements
    testing individual flags. The functionality for handling each flag is
    unchanged.

    Link: https://lore.kernel.org/r/1622827263-12516-3-git-send-email-mikelley@microsoft.com
    Signed-off-by: Michael Kelley <mikelley@microsoft.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0c5085c93f641a9def4e790a18d09749021b3066
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Thu May 20 13:50:38 2021 +0000

    powerpc/inst: Fix sparse detection on get_user_instr()

    [ Upstream commit b3a9e523237013477bea914b7fbfbe420428b988 ]

    get_user_instr() lacks sparse detection for the __user tag.

    This is because __gui_ptr is assigned with a cast.

    Fix that by adding a __chk_user_ptr()

    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/0320e5b41a794fd456ab8c5993bbfadcf9e1d8b4.1621516826.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 71bfeb9e0f1ac09138017f107d0350db5302c949
Author: Gil Fine <gil.fine@intel.com>
Date:   Mon Jun 14 16:52:10 2021 +0300

    thunderbolt: Fix DROM handling for USB4 DROM

    [ Upstream commit b18f901382fdb74a138a0bf30458c54a023a1d86 ]

    DROM for USB4 host/device has a shorter header than Thunderbolt DROM
    header. This patch addresses host/device with USB4 DROM (According to spec:
    Universal Serial Bus 4 (USB4) Device ROM Specification, Rev 1.0, Feb-2021).

    While there correct the data_len field to be 12 bits and rename
    __unknown1 to reserved following the spec.

    Signed-off-by: Gil Fine <gil.fine@intel.com>
    Signed-off-by: Mika Westerberg <mika.westerberg@linux.intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d18303a5bc114d5b50f1e1d6ebf2c23e22f9dbc
Author: Srinivas Neeli <srinivas.neeli@xilinx.com>
Date:   Fri Apr 9 19:38:06 2021 +0530
```

```
      gpio: zynq: Check return value of irq_get_irq_data

      [ Upstream commit 35d7b72a632bc7afb15356f44005554af8697904 ]

      In two different instances the return value of "irq_get_irq_data"
      API was neither captured nor checked.
      Fixed it by capturing the return value and then checking for any error.

      Addresses-Coverity: "returned_null"
      Signed-off-by: Srinivas Neeli <srinivas.neeli@xilinx.com>
      Signed-off-by: Bartosz Golaszewski <bgolaszewski@baylibre.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c42ec0450a55f74871f14667cb356ed4e45a8d4e
  Author: Srinivas Neeli <srinivas.neeli@xilinx.com>
  Date:   Fri Apr 9 19:38:05 2021 +0530

      gpio: zynq: Check return value of pm_runtime_get_sync

      [ Upstream commit a51b2fb94b04ab71e53a71b9fad03fa826941254 ]

      Return value of "pm_runtime_get_sync" API was neither captured nor checked.
      Fixed it by capturing the return value and then checking for any warning.

      Addresses-Coverity: "check_return"
      Signed-off-by: Srinivas Neeli <srinivas.neeli@xilinx.com>
      Signed-off-by: Bartosz Golaszewski <bgolaszewski@baylibre.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b5b9d977c7a8c7bc18dc538923ecc1c8a0238298
  Author: Jaroslav Kysela <perex@perex.cz>
  Date:   Mon Jun 14 09:17:46 2021 +0200

      ASoC: soc-pcm: fix the return value in dpcm_apply_symmetry()

      [ Upstream commit 12ffd726824a2f52486f72338b6fd3244b512959 ]

      In case, where the loops are not executed for a reason, the uninitialized
      variable 'err' is returned to the caller. Make code fully predictible
      and assign zero in the declaration.

      Signed-off-by: Jaroslav Kysela <perex@perex.cz>
      Cc: Mark Brown <broonie@kernel.org>
      Cc: Kuninori Morimoto <kuninori.morimoto.gx@renesas.com>
      Link: https://lore.kernel.org/r/20210614071746.1787072-1-perex@perex.cz
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 77921518bf652e3c894182501f7380c6eb992cc1
  Author: Jaroslav Kysela <perex@perex.cz>
  Date:   Mon Jun 14 09:17:10 2021 +0200

      ALSA: control_led - fix initialization in the mode show callback

      [ Upstream commit e381a14c3e3a4e90e293d4eaa5a3ab8ae98b9973 ]

      The str variable should be always initialized before use even if
      the switch covers all cases. This is a minimalistic fix: Assign NULL,
      the sprintf() may print '(null)' if something is corrupted.

      Signed-off-by: Jaroslav Kysela <perex@perex.cz>
      Link: https://lore.kernel.org/r/20210614071710.1786866-1-perex@perex.cz
      Signed-off-by: Takashi Iwai <tiwai@suse.de>
      Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 2955c54c5bca08aa10cb5c8704130eb10d4937d5
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Thu Jun 10 20:49:58 2021 +0800

    ALSA: n64: check return value after calling platform_get_resource()

    [ Upstream commit be471fe332f7f14aa6828010b220d7e6902b91a0 ]

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

    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
    Link: https://lore.kernel.org/r/20210610124958.116142-1-yangyingliang@huawei.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f9741a9a91f25c89e04b408cd61e3ab050ce24b
Author: Xiyu Yang <xiyuyang19@fudan.edu.cn>
Date:   Thu Jun 10 10:49:20 2021 +0800

    iommu/arm-smmu: Fix arm_smmu_device refcount leak in address translation

    [ Upstream commit 7c8f176d6a3fa18aa0f8875da6f7c672ed2a8554 ]

    The reference counting issue happens in several exception handling paths
    of arm_smmu_iova_to_phys_hard(). When those error scenarios occur, the
    function forgets to decrease the refcount of "smmu" increased by
    arm_smmu_rpm_get(), causing a refcount leak.

    Fix this issue by jumping to "out" label when those error scenarios
    occur.

    Signed-off-by: Xiyu Yang <xiyuyang19@fudan.edu.cn>
    Signed-off-by: Xin Tan <tanxin.ctf@gmail.com>
    Reviewed-by: Rob Clark <robdclark@chromium.org>
    Link: https://lore.kernel.org/r/1623293391-17261-1-git-send-email-
xiyuyang19@fudan.edu.cn
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fe92c058199067ae90cf2a901ddf3c271893557a
Author: Xiyu Yang <xiyuyang19@fudan.edu.cn>
Date:   Thu Jun 10 10:54:29 2021 +0800

    iommu/arm-smmu: Fix arm_smmu_device refcount leak when arm_smmu_rpm_get fails

    [ Upstream commit 1adf30f198c26539a62d761e45af72cde570413d ]

    arm_smmu_rpm_get() invokes pm_runtime_get_sync(), which increases the
    refcount of the "smmu" even though the return value is less than 0.

    The reference counting issue happens in some error handling paths of
    arm_smmu_rpm_get() in its caller functions. When arm_smmu_rpm_get()
    fails, the caller functions forget to decrease the refcount of "smmu"
    increased by arm_smmu_rpm_get(), causing a refcount leak.

    Fix this issue by calling pm_runtime_resume_and_get() instead of
    pm_runtime_get_sync() in arm_smmu_rpm_get(), which can keep the refcount
    balanced in case of failure.

    Signed-off-by: Xiyu Yang <xiyuyang19@fudan.edu.cn>
    Signed-off-by: Xin Tan <tanxin.ctf@gmail.com>
    Link: https://lore.kernel.org/r/1623293672-17954-1-git-send-email-
xiyuyang19@fudan.edu.cn
    Signed-off-by: Will Deacon <will@kernel.org>
```

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

commit f30d6b6bb4f6ad98d04d4aa7b6c9236de557e33f
Author: Geoff Levand <geoff@infradead.org>
Date:   Thu Jun 3 19:17:02 2021 +0000

    powerpc/ps3: Add dma_mask to ps3_dma_region

    [ Upstream commit 9733862e50fdba55e7f1554e4286fcc5302ff28e ]

    Commit f959dcd6ddfd29235030e8026471ac1b022ad2b0 (dma-direct: Fix
    potential NULL pointer dereference) added a null check on the
    dma_mask pointer of the kernel's device structure.

    Add a dma_mask variable to the ps3_dma_region structure and set
    the device structure's dma_mask pointer to point to this new variable.

    Fixes runtime errors like these:
    # WARNING: Fixes tag on line 10 doesn't match correct format
    # WARNING: Fixes tag on line 10 doesn't match correct format

      ps3_system_bus_match:349: dev=8.0(sb_01), drv=8.0(ps3flash): match
      WARNING: CPU: 0 PID: 1 at kernel/dma/mapping.c:151 .dma_map_page_attrs+0x34/0x1e0
      ps3flash sb_01: ps3stor_setup:193: map DMA region failed

    Signed-off-by: Geoff Levand <geoff@infradead.org>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://lore.kernel.org/r/562d0c9ea0100a30c3b186bcc7adb34b0bbd2cd7.1622746428.git.geoff@infr
adead.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 06771032116a5e9c923b0b53294b6552eab6398c
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Jun 8 16:04:37 2021 +0200

    ALSA: sb: Fix potential double-free of CSP mixer elements

    [ Upstream commit c305366a37441c2ac90b08711cb6f032b43672f2 ]

    snd_sb_qsound_destroy() contains the calls of removing the previously
    created mixer controls, but it doesn't clear the pointers.  As
    snd_sb_qsound_destroy() itself may be repeatedly called via ioctl,
    this could lead to double-free potentially.

    Fix it by clearing the struct fields properly afterwards.

    Link: https://lore.kernel.org/r/20210608140540.17885-4-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c33fbdc1cdaefe324c2f33304245ca42f251e1ed
Author: Eric Anholt <eric@anholt.net>
Date:   Fri Mar 26 16:13:02 2021 -0700

    iommu/arm-smmu-qcom: Skip the TTBR1 quirk for db820c.

    [ Upstream commit a242f4297cfe3f4589a7620dcd42cc503607fc6b ]

    db820c wants to use the qcom smmu path to get HUPCF set (which keeps
    the GPU from wedging and then sometimes wedging the kernel after a
    page fault), but it doesn't have separate pagetables support yet in
    drm/msm so we can't go all the way to the TTBR1 path.

    Signed-off-by: Eric Anholt <eric@anholt.net>
```

```
    Reviewed-by: Bjorn Andersson <bjorn.andersson@linaro.org>
    Link: https://lore.kernel.org/r/20210326231303.3071950-1-eric@anholt.net
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de544803cbe02ab2acb9c1040e155ea3a38b3d95
Author: Po-Hsu Lin <po-hsu.lin@canonical.com>
Date:   Tue May 25 10:36:14 2021 +0800

    selftests: timers: rtcpie: skip test if default RTC device does not exist

    [ Upstream commit 0d3e5a057992bdc66e4dca2ca50b77fa4a7bd90e ]

    This test will require /dev/rtc0, the default RTC device, or one
    specified by user to run. Since this default RTC is not guaranteed to
    exist on all of the devices, so check its existence first, otherwise
    skip this test with the kselftest skip code 4.

    Without this patch this test will fail like this on a s390x zVM:
    $ selftests: timers: rtcpie
    $ /dev/rtc0: No such file or directory
    not ok 1 selftests: timers: rtcpie # exit=22

    With this patch:
    $ selftests: timers: rtcpie
    $ Default RTC /dev/rtc0 does not exist. Test Skipped!
    not ok 9 selftests: timers: rtcpie # SKIP

    Fixed up change log so "With this patch" text doesn't get dropped.
    Shuah Khan <skhan@linuxfoundation.org>

    Signed-off-by: Po-Hsu Lin <po-hsu.lin@canonical.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 52e8bd03856a019667f5a449f8f4a015d5a80d56
Author: Fabrice Fontaine <fontaine.fabrice@gmail.com>
Date:   Mon May 10 07:31:33 2021 +0200

    s390: disable SSP when needed

    [ Upstream commit 42e8d652438f5ddf04e5dac299cb5e623d113dc0 ]

    Though -nostdlib is passed in PURGATORY_LDFLAGS and -ffreestanding in
    KBUILD_CFLAGS_DECOMPRESSOR, -fno-stack-protector must also be passed to
    avoid linking errors related to undefined references to
    '__stack_chk_guard' and '__stack_chk_fail' if toolchain enforces
    -fstack-protector.

    Fixes
     - https://gitlab.com/kubu93/buildroot/-/jobs/1247043361

    Signed-off-by: Fabrice Fontaine <fontaine.fabrice@gmail.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Reviewed-by: Alexander Egorenkov <egorenar@linux.ibm.com>
    Tested-by: Alexander Egorenkov <egorenar@linux.ibm.com>
    Link: https://lore.kernel.org/r/20210510053133.1220167-1-fontaine.fabrice@gmail.com
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2fdbf6e7baa86917700a73aa334e502efb62219b
Author: Valentin Vidic <vvidic@valentin-vidic.from.hr>
Date:   Tue Apr 27 21:40:10 2021 +0200

    s390/sclp_vt220: fix console name to match device
```

```
    [ Upstream commit b7d91d230a119fdcc334d10c9889ce9c5e15118b ]

    Console name reported in /proc/consoles:

      ttyS1                -W- (EC p  )    4:65

    does not match the char device name:

      crw--w----    1 root     root        4,  65 May 17 12:18 /dev/ttysclp0

    so debian-installer inside a QEMU s390x instance gets confused and fails
    to start with the following error:

      steal-ctty: No such file or directory

    Signed-off-by: Valentin Vidic <vvidic@valentin-vidic.from.hr>
    Link: https://lore.kernel.org/r/20210427194010.9330-1-vvidic@valentin-vidic.from.hr
    Signed-off-by: Christian Borntraeger <borntraeger@de.ibm.com>
    Signed-off-by: Vasily Gorbik <gor@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b29ddd99c14fb494739cdf3a6a7658616eae63a7
Author: Daniel Mack <daniel@zonque.org>
Date:   Fri May 28 15:33:20 2021 +0200

    serial: tty: uartlite: fix console setup

    [ Upstream commit d157fca711ad42e75efef3444c83d2e1a17be27a ]

    Remove the hack to assign the global console_port variable at probe time.
    This assumption that cons->index is -1 is wrong for systems that specify
    'console=' in the cmdline (or 'stdout-path' in dts). Hence, on such system
    the actual console assignment is ignored, and the first UART that happens
    to be probed is used as console instead.

    Move the logic to console_setup() and map the console to the correct port
    through the array of available ports instead.

    Signed-off-by: Daniel Mack <daniel@zonque.org>
    Link: https://lore.kernel.org/r/20210528133321.1859346-1-daniel@zonque.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2124240be810f7a13a7524269a2d92bf3cc5c362
Author: Zou Wei <zou_wei@huawei.com>
Date:   Thu May 13 16:57:29 2021 +0800

    fsi: Add missing MODULE_DEVICE_TABLE

    [ Upstream commit 19a52178125c1e8b84444d85f2ce34c0964b4a91 ]

    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/1620896249-52769-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Joel Stanley <joel@jms.id.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed8c2cc773811a18f84f9f89e4662183d28177f6
Author: Al Viro <viro@zeniv.linux.org.uk>
Date:   Fri Apr 23 22:24:08 2021 -0400
```

```
    iov_iter_advance(): use consistent semantics for move past the end

    [ Upstream commit 3b3fc051cd2cba42bf736fa62780857d251a1236 ]

    asking to advance by more than we have left in the iov_iter should
    move to the very end; it should *not* leave negative i->count and
    it should not spew into syslog, etc. - it's a legitimate operation.

    Signed-off-by: Al Viro <viro@zeniv.linux.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3a117833731606baaed25d39aaae70a4a501221
Author: Yufen Yu <yuyufen@huawei.com>
Date:   Mon May 24 05:35:21 2021 -0400

    ASoC: img: Fix PM reference leak in img_i2s_in_probe()

    [ Upstream commit 81aad47278539f02de808bcc8251fed0ad3d6f55 ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yufen Yu <yuyufen@huawei.com>
    Link: https://lore.kernel.org/r/20210524093521.612176-1-yuyufen@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc78632eb0857e56ac0e228266a050879da3a139
Author: Tony Lindgren <tony@atomide.com>
Date:   Sat May 22 08:10:01 2021 +0300

    mfd: cpcap: Fix cpcap dmamask not set warnings

    [ Upstream commit 0b7cbe811ca524295ea43d9a4d73d3427e419c54 ]

    We have started to get a bunch of pointless dmamask not set warnings
    that makes the output of dmesg -l err,warn hard to read with many
    extra warnings:

    cpcap-regulator cpcap-regulator.0: DMA mask not set
    cpcap_adc cpcap_adc.0: DMA mask not set
    cpcap_battery cpcap_battery.0: DMA mask not set
    cpcap-charger cpcap-charger.0: DMA mask not set
    cpcap-pwrbutton cpcap-pwrbutton.0: DMA mask not set
    cpcap-led cpcap-led.0: DMA mask not set
    cpcap-led cpcap-led.1: DMA mask not set
    cpcap-led cpcap-led.2: DMA mask not set
    cpcap-led cpcap-led.3: DMA mask not set
    cpcap-led cpcap-led.4: DMA mask not set
    cpcap-rtc cpcap-rtc.0: DMA mask not set
    cpcap-usb-phy cpcap-usb-phy.0: DMA mask not set

    This seems to have started with commit 4d8bde883bfb ("OF: Don't set
    default coherent DMA mask"). We have the parent SPI controller use
    DMA, while CPCAP driver and it's children do not. For audio, the
    DMA is handled over I2S bus with the McBSP driver.

    Cc: Carl Philipp Klemm <philipp@uvos.xyz>
    Cc: Ivan Jelincic <parazyd@dyne.org>
    Cc: Merlijn Wajer <merlijn@wizzup.org>
    Cc: Pavel Machek <pavel@ucw.cz>
```

```
    Cc: Sebastian Reichel <sre@kernel.org>
    Cc: Sicelo A. Mhlongo <absicsz@gmail.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4aee78c62b9d04d746bec4d85e8a2f8e5c5cd217
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 14:33:46 2021 +0800

    mfd: da9052/stmpe: Add and modify MODULE_DEVICE_TABLE

    [ Upstream commit 4700ef326556ed74aba188f12396740a8c1c21dd ]

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

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Zou Wei <zou_wei@huawei.com>
    Signed-off-by: Lee Jones <lee.jones@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e65432dce6b5cb1daf6864f900440c78d0e51258
Author: Mike Christie <michael.christie@oracle.com>
Date:   Tue May 25 13:18:18 2021 -0500

    scsi: qedi: Fix cleanup session block/unblock use

    [ Upstream commit 0c72191da68638a479602dd515b587ada913184a ]

    Drivers shouldn't be calling block/unblock session for cmd cleanup because
    the functions can change the session state from under libiscsi.  This adds
    a new a driver level bit so it can block all I/O the host while it drains
    the card.

    Link: https://lore.kernel.org/r/20210525181821.7617-26-michael.christie@oracle.com
    Reviewed-by: Manish Rangankar <mrangankar@marvell.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 400c206f3be09d8aa23cbb230c985875d35b9105
Author: Mike Christie <michael.christie@oracle.com>
Date:   Tue May 25 13:18:17 2021 -0500

    scsi: qedi: Fix TMF session block/unblock use

    [ Upstream commit 2819b4ae2873d50fd55292877b0231ec936c3b2e ]

    Drivers shouldn't be calling block/unblock session for tmf handling because
    the functions can change the session state from under libiscsi.
    iscsi_queuecommand's call to iscsi_prep_scsi_cmd_pdu->
    iscsi_check_tmf_restrictions will prevent new cmds from being sent to qedi
    after we've started handling a TMF. So we don't need to try and block it in
    the driver, and we can remove these block calls.

    Link: https://lore.kernel.org/r/20210525181821.7617-25-michael.christie@oracle.com
    Reviewed-by: Manish Rangankar <mrangankar@marvell.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf5cc57ea7dc817e1f00b35164aa9dfa08685066
Author: Mike Christie <michael.christie@oracle.com>
```

```
Date:   Tue May 25 13:18:13 2021 -0500

    scsi: qedi: Fix race during abort timeouts

    [ Upstream commit 2ce002366a3fcc3f9616d4583194f65dde0ad253 ]

    If the SCSI cmd completes after qedi_tmf_work calls iscsi_itt_to_task then
    the qedi qedi_cmd->task_id could be freed and used for another cmd. If we
    then call qedi_iscsi_cleanup_task with that task_id we will be cleaning up
    the wrong cmd.

    Wait to release the task_id until the last put has been done on the
    iscsi_task. Because libiscsi grabs a ref to the task when sending the
    abort, we know that for the non-abort timeout case that the task_id we are
    referencing is for the cmd that was supposed to be aborted.

    A latter commit will fix the case where the abort times out while we are
    running qedi_tmf_work.

    Link: https://lore.kernel.org/r/20210525181821.7617-21-michael.christie@oracle.com
    Reviewed-by: Manish Rangankar <mrangankar@marvell.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 338330351ca52fce2ad1bd6d1cafc5bb426edbd0
Author: Mike Christie <michael.christie@oracle.com>
Date:   Tue May 25 13:18:12 2021 -0500

    scsi: qedi: Fix null ref during abort handling

    [ Upstream commit 5777b7f0f03ce49372203b6521631f62f2810c8f ]

    If qedi_process_cmd_cleanup_resp finds the cmd it frees the work and sets
    list_tmf_work to NULL, so qedi_tmf_work should check if list_tmf_work is
    non-NULL when it wants to force cleanup.

    Link: https://lore.kernel.org/r/20210525181821.7617-20-michael.christie@oracle.com
    Reviewed-by: Manish Rangankar <mrangankar@marvell.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 52b5bf245a682f9013c1b6ebddaff7aadd6f5e34
Author: Mike Christie <michael.christie@oracle.com>
Date:   Tue May 25 13:18:07 2021 -0500

    scsi: iscsi: Fix shost->max_id use

    [ Upstream commit bdd4aad7ff92ae39c2e93c415bb6761cb8b584da ]

    The iscsi offload drivers are setting the shost->max_id to the max number
    of sessions they support. The problem is that max_id is not the max number
    of targets but the highest identifier the targets can have. To use it to
    limit the number of targets we need to set it to max sessions - 1, or we
    can end up with a session we might not have preallocated resources for.

    Link: https://lore.kernel.org/r/20210525181821.7617-15-michael.christie@oracle.com
    Reviewed-by: Lee Duncan <lduncan@suse.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fa9542b35ceb4202e8f8d65f440529a63524dca9
Author: Mike Christie <michael.christie@oracle.com>
```

```
Date:   Tue May 25 13:18:06 2021 -0500

    scsi: iscsi: Fix conn use after free during resets

    [ Upstream commit ec29d0ac29be366450a7faffbcf8cba3a6a3b506 ]

    If we haven't done a unbind target call we can race where
    iscsi_conn_teardown wakes up the EH thread and then frees the conn while
    those threads are still accessing the conn ehwait.

    We can only do one TMF per session so this just moves the TMF fields from
    the conn to the session. We can then rely on the
    iscsi_session_teardown->iscsi_remove_session->__iscsi_unbind_session call
    to remove the target and it's devices, and know after that point there is
    no device or scsi-ml callout trying to access the session.

    Link: https://lore.kernel.org/r/20210525181821.7617-14-michael.christie@oracle.com
    Reviewed-by: Lee Duncan <lduncan@suse.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf3ce567db5639aab2010a49b4e522830e0887a0
Author: Mike Christie <michael.christie@oracle.com>
Date:   Tue May 25 13:18:03 2021 -0500

    scsi: iscsi: Add iscsi_cls_conn refcount helpers

    [ Upstream commit b1d19e8c92cfb0ded180ef3376c20e130414e067 ]

    There are a couple places where we could free the iscsi_cls_conn while it's
    still in use. This adds some helpers to get/put a refcount on the struct
    and converts an exiting user. Subsequent commits will then use the helpers
    to fix 2 bugs in the eh code.

    Link: https://lore.kernel.org/r/20210525181821.7617-11-michael.christie@oracle.com
    Reviewed-by: Lee Duncan <lduncan@suse.com>
    Signed-off-by: Mike Christie <michael.christie@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1420af019bcb9d247dc9de2335c4e6b48144ce6d
Author: Chandrakanth Patil <chandrakanth.patil@broadcom.com>
Date:   Fri May 28 18:43:06 2021 +0530

    scsi: megaraid_sas: Handle missing interrupts while re-enabling IRQs

    [ Upstream commit 9bedd36e9146b34dda4d6994e3aa1d72bc6442c1 ]

    While reenabling the IRQ after IRQ poll there may be a small window for the
    firmware to post the replies with interrupts raised. In that case the
    driver will not see the interrupts which leads to I/O timeout.

    This issue only happens when there are many I/O completions on a single
    reply queue. This forces the driver to switch between the interrupt and IRQ
    context.

    Make the driver process the reply queue one more time after enabling the
    IRQ.

    Link: https://lore.kernel.org/linux-scsi/20201102072746.27410-1-
sreekanth.reddy@broadcom.com/
    Link: https://lore.kernel.org/r/20210528131307.25683-5-chandrakanth.patil@broadcom.com
    Cc: Tomas Henzl <thenzl@redhat.com>
    Reported-by: kernel test robot <lkp@intel.com>
```

```
    Signed-off-by: Chandrakanth Patil <chandrakanth.patil@broadcom.com>
    Signed-off-by: Sumit Saxena <sumit.saxena@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2262bb7ee19e2d0056ffae84cd1803bd330d06af
Author: Kashyap Desai <kashyap.desai@broadcom.com>
Date:   Fri May 28 18:43:05 2021 +0530

    scsi: megaraid_sas: Early detection of VD deletion through RaidMap update

    [ Upstream commit ae6874ba4b43c5a00065f48599811a09d33b873d ]

    Consider the case where a VD is deleted and the targetID of that VD is
    assigned to a newly created VD. If the sequence of deletion/addition of VD
    happens very quickly there is a possibility that second event (VD add)
    occurs even before the driver processes the first event (VD delete).  As
    event processing is done in deferred context the device list remains the
    same (but targetID is re-used) so driver will not learn the VD
    deletion/additon. I/Os meant for the older VD will be directed to new VD
    which may lead to data corruption.

    Make driver detect the deleted VD as soon as possible based on the RaidMap
    update and block further I/O to that device.

    Link: https://lore.kernel.org/r/20210528131307.25683-4-chandrakanth.patil@broadcom.com
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Kashyap Desai <kashyap.desai@broadcom.com>
    Signed-off-by: Chandrakanth Patil <chandrakanth.patil@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e623f79691c5104317669ab36ec316a90c05062f
Author: Chandrakanth Patil <chandrakanth.patil@broadcom.com>
Date:   Fri May 28 18:43:04 2021 +0530

    scsi: megaraid_sas: Fix resource leak in case of probe failure

    [ Upstream commit b5438f48fdd8e1c3f130d32637511efd32038152 ]

    The driver doesn't clean up all the allocated resources properly when
    scsi_add_host(), megasas_start_aen() function fails during the PCI device
    probe.

    Clean up all those resources.

    Link: https://lore.kernel.org/r/20210528131307.25683-3-chandrakanth.patil@broadcom.com
    Signed-off-by: Chandrakanth Patil <chandrakanth.patil@broadcom.com>
    Signed-off-by: Sumit Saxena <sumit.saxena@broadcom.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0dab71667eb029780333ff45b07e3fcce5e16119
Author: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
Date:   Tue Jun 1 19:07:10 2021 +0800

    fs/jfs: Fix missing error code in lmLogInit()

    [ Upstream commit 492109333c29e1bb16d8732e1d597b02e8e0bf2e ]

    The error code is missing in this code scenario, add the error code
    '-EINVAL' to the return value 'rc.

    Eliminate the follow smatch warning:
```

```
      fs/jfs/jfs_logmgr.c:1327 lmLogInit() warn: missing error code 'rc'.

      Reported-by: Abaci Robot <abaci@linux.alibaba.com>
      Signed-off-by: Jiapeng Chong <jiapeng.chong@linux.alibaba.com>
      Signed-off-by: Dave Kleikamp <dave.kleikamp@oracle.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 936440f4782692b0ee7fb284bb9f302d22c527ce
  Author: Hannes Reinecke <hare@suse.de>
  Date:   Tue Apr 27 10:30:11 2021 +0200

      scsi: scsi_dh_alua: Check for negative result value

      [ Upstream commit 7e26e3ea028740f934477ec01ba586ab033c35aa ]

      scsi_execute() will now return a negative error if there was an error prior
      to command submission; evaluate that instead if checking for DRIVER_ERROR.

      [mkp: build fix]

      Link: https://lore.kernel.org/r/20210427083046.31620-6-hare@suse.de
      Signed-off-by: Hannes Reinecke <hare@suse.de>
      Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 89b033dd8f3b537f664702fca8c95a42a2ddf1a5
  Author: Hannes Reinecke <hare@suse.de>
  Date:   Tue Apr 27 10:30:09 2021 +0200

      scsi: core: Fixup calling convention for scsi_mode_sense()

      [ Upstream commit 8793613de913e03e7c884f4cc56e350bc716431e ]

      The description for scsi_mode_sense() claims to return the number of valid
      bytes on success, which is not what the code does.  Additionally there is
      no gain in returning the SCSI status, as everything the callers do is to
      check against scsi_result_is_good(), which is what scsi_mode_sense() does
      already.  So change the calling convention to return a standard error code
      on failure, and 0 on success, and adapt the description and all callers.

      Link: https://lore.kernel.org/r/20210427083046.31620-4-hare@suse.de
      Reviewed-by: Bart Van Assche <bvanassche@acm.org>
      Signed-off-by: Hannes Reinecke <hare@suse.de>
      Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fc6bf9941021d6baad216a9d5831450f3fa9ccd7
  Author: Suganath Prabu S <suganath-prabu.subramani@broadcom.com>
  Date:   Tue May 18 10:46:23 2021 +0530

      scsi: mpt3sas: Fix deadlock while cancelling the running firmware event

      [ Upstream commit e2fac6c44ae06e58ac02181b048af31195883c31 ]

      Do not cancel current running firmware event work if the event type is
      different from MPT3SAS_REMOVE_UNRESPONDING_DEVICES.  Otherwise a deadlock
      can be observed while cancelling the current firmware event work if a hard
      reset operation is called as part of processing the current event.

      Link: https://lore.kernel.org/r/20210518051625.1596742-2-suganath-
  prabu.subramani@broadcom.com
      Signed-off-by: Suganath Prabu S <suganath-prabu.subramani@broadcom.com>
      Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit b2ef1f5de40342de44fc5355321595f91774dab5
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Tue May 25 21:44:04 2021 +0200

    tty: serial: 8250: serial_cs: Fix a memory leak in error handling path

    [ Upstream commit fad92b11047a748c996ebd6cfb164a63814eeb2e ]

    In the probe function, if the final 'serial_config()' fails, 'info' is
    leaking.

    Add a resource handling path to free this memory.

    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Link:
https://lore.kernel.org/r/dc25f96b7faebf42e60fe8d02963c941cf4d8124.1621971720.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b731dddd686b1a5cd0c42c54b47439375fb83b15
Author: Lucas Tanure <tanureal@opensource.cirrus.com>
Date:   Tue May 25 10:08:19 2021 +0100

    ASoC: cs42l42: Fix 1536000 Bit Clock instability

    [ Upstream commit 1c52825c38fc4e44c61ed75a8ae32f5fa580383b ]

    The 16 Bits, 2 channels, 48K sample rate use case needs
    to configure a safer pll_divout during the start of PLL
    After 800us from the start of PLL the correct pll_divout
    can be set

    Signed-off-by: Lucas Tanure <tanureal@opensource.cirrus.com>
    Reviewed-by: Richard Fitzgerald <rf@opensource.cirrus.com>
    Message-Id: <20210525090822.64577-1-tanureal@opensource.cirrus.com>
    Signed-off-by: Mark Brown <broonie@sirena.org.uk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2a1115ad5fce313e90e1402176fe1fd1296e6aa9
Author: Yufen Yu <yuyufen@huawei.com>
Date:   Mon May 24 05:38:11 2021 -0400

    ALSA: ac97: fix PM reference leak in ac97_bus_remove()

    [ Upstream commit a38e93302ee25b2ca6f4ee76c6c974cf3637985e ]

    pm_runtime_get_sync will increment pm usage counter even it failed.
    Forgetting to putting operation will result in reference leak here.
    Fix it by replacing it with pm_runtime_resume_and_get to keep usage
    counter balanced.

    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Yufen Yu <yuyufen@huawei.com>
    Link: https://lore.kernel.org/r/20210524093811.612302-1-yuyufen@huawei.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dcb634e434b850287b31a117637cbf598563ce1c
Author: John Garry <john.garry@huawei.com>
Date:   Wed May 19 22:31:02 2021 +0800

    scsi: core: Cap scsi_host cmd_per_lun at can_queue

    [ Upstream commit ea2f0f77538c50739b9fb4de4700cee5535e1f77 ]
```

```
    The sysfs handling function sdev_store_queue_depth() enforces that the sdev
    queue depth cannot exceed shost can_queue. The initial sdev queue depth
    comes from shost cmd_per_lun. However, the LLDD may manually set
    cmd_per_lun to be larger than can_queue, which leads to an initial sdev
    queue depth greater than can_queue.

    Such an issue was reported in [0], which caused a hang. That has since been
    fixed in commit fc09acb7de31 ("scsi: scsi_debug: Fix cmd_per_lun, set to
    max_queue").

    Stop this possibly happening for other drivers by capping shost cmd_per_lun
    at shost can_queue.

    [0] https://lore.kernel.org/linux-scsi/YHaez6iN2HHYxYOh@T590/

    Link: https://lore.kernel.org/r/1621434662-173079-1-git-send-email-john.garry@huawei.com
    Reviewed-by: Ming Lei <ming.lei@redhat.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b927240f44c62d490770c744eb09d2b4a572191
Author: James Smart <jsmart2021@gmail.com>
Date:   Fri May 14 12:55:56 2021 -0700

    scsi: lpfc: Fix crash when lpfc_sli4_hba_setup() fails to initialize the SGLs

    [ Upstream commit 5aa615d195f1e142c662cb2253f057c9baec7531 ]

    The driver is encountering a crash in lpfc_free_iocb_list() while
    performing initial attachment.

    Code review found this to be an errant failure path that was taken, jumping
    to a tag that then referenced structures that were uninitialized.

    Fix the failure path.

    Link: https://lore.kernel.org/r/20210514195559.119853-9-jsmart2021@gmail.com
    Co-developed-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: James Smart <jsmart2021@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c1a1c3e71dd7da82fa9a549105263f2088f1ad3
Author: James Smart <jsmart2021@gmail.com>
Date:   Fri May 14 12:55:51 2021 -0700

    scsi: lpfc: Fix "Unexpected timeout" error in direct attach topology

    [ Upstream commit e30d55137edef47434c40d7570276a0846fe922c ]

    An 'unexpected timeout' message may be seen in a point-2-point topology.
    The message occurs when a PLOGI is received before the driver is notified
    of FLOGI completion. The FLOGI completion failure causes discovery to be
    triggered for a second time. The discovery timer is restarted but no new
    discovery activity is initiated, thus the timeout message eventually
    appears.

    In point-2-point, when discovery has progressed before the FLOGI completion
    is processed, it is not a failure. Add code to FLOGI completion to detect
    that discovery has progressed and exit the FLOGI handling (noop'ing it).
```

```
    Link: https://lore.kernel.org/r/20210514195559.119853-4-jsmart2021@gmail.com
    Co-developed-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: Justin Tee <justin.tee@broadcom.com>
    Signed-off-by: James Smart <jsmart2021@gmail.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd3bc6dc69c6d01c05382b94844fdfde91f229a6
Author: Sergey Shtylyov <s.shtylyov@omp.ru>
Date:   Wed May 19 22:20:15 2021 +0300

    scsi: hisi_sas: Propagate errors in interrupt_init_v1_hw()

    [ Upstream commit ab17122e758ef68fb21033e25c041144067975f5 ]

    After commit 6c11dc060427 ("scsi: hisi_sas: Fix IRQ checks") we have the
    error codes returned by platform_get_irq() ready for the propagation
    upsream in interrupt_init_v1_hw() -- that will fix still broken deferred
    probing. Let's propagate the error codes from devm_request_irq() as well
    since I don't see the reason to override them with -ENOENT...

    Link: https://lore.kernel.org/r/49ba93a3-d427-7542-d85a-b74fe1a33a73@omp.ru
    Acked-by: John Garry <john.garry@huawei.com>
    Signed-off-by: Sergey Shtylyov <s.shtylyov@omp.ru>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 67fb59eceee3fbd8dbab34ab9a6d53983435fdb4
Author: ching Huang <ching2048@areca.com.tw>
Date:   Thu May 20 14:55:15 2021 +0800

    scsi: arcmsr: Fix doorbell status being updated late on ARC-1886

    [ Upstream commit d9a231226f28261a787535e08d0c78669e1ad010 ]

    It is possible for the IOP to be delayed in updating the doorbell
    status. The doorbell status should not be 0 so loop until the value
    changes.

    Link:
https://lore.kernel.org/r/afdfdf7eabecf14632492c4987a6b9ac6312a7ad.camel@areca.com.tw
    Signed-off-by: ching Huang <ching2048@areca.com.tw>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab935078f3613cd042166d6235530def418b3285
Author: Luiz Sampaio <sampaio.ime@gmail.com>
Date:   Wed May 19 19:30:44 2021 -0300

    w1: ds2438: fixing bug that would always get page0

    [ Upstream commit 1f5e7518f063728aee0679c5086b92d8ea429e11 ]

    The purpose of the w1_ds2438_get_page function is to get the register
    values at the page passed as the pageno parameter. However, the page0 was
    hardcoded, such that the function always returned the page0 contents. Fixed
    so that the function can retrieve any page.

    Signed-off-by: Luiz Sampaio <sampaio.ime@gmail.com>
    Link: https://lore.kernel.org/r/20210519223046.13798-5-sampaio.ime@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 59f03af920b571ea0333a6d542d3ce222026926d
Author: Jaska Uimonen <jaska.uimonen@linux.intel.com>
```

```
Date:   Fri May 21 12:28:02 2021 +0300

    ASoC: SOF: topology: fix assignment to use le32_to_cpu

    [ Upstream commit ccaea61a8d1b8180cc3c470e383381884e4bc1f2 ]

    Fix sparse warning by using le32_to_cpu.

    Signed-off-by: Jaska Uimonen <jaska.uimonen@linux.intel.com>
    Reviewed-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Ranjani Sridharan <ranjani.sridharan@linux.intel.com>
    Signed-off-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
    Link: https://lore.kernel.org/r/20210521092804.3721324-6-kai.vehmanen@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a133a0996d6b4c83509d570ed4edcba34c44f25
Author: Chunfeng Yun <chunfeng.yun@mediatek.com>
Date:   Wed May 19 14:39:44 2021 +0800

    usb: common: usb-conn-gpio: fix NULL pointer dereference of charger

    [ Upstream commit 880287910b1892ed2cb38977893b947382a09d21 ]

    When power on system with OTG cable, IDDIG's interrupt arises before
    the charger registration, it will cause a NULL pointer dereference,
    fix the issue by registering the power supply before requesting
    IDDIG/VBUS irq.

    Signed-off-by: Chunfeng Yun <chunfeng.yun@mediatek.com>
    Link: https://lore.kernel.org/r/1621406386-18838-1-git-send-email-
chunfeng.yun@mediatek.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b57a2dd956c8a0dc0a824cc1baeb62443513ef7
Author: Takashi Sakamoto <o-takashi@sakamocchi.jp>
Date:   Tue May 18 17:45:47 2021 +0900

    Revert "ALSA: bebob/oxfw: fix Kconfig entry for Mackie d.2 Pro"

    [ Upstream commit 5d6fb80a142b5994355ce675c517baba6089d199 ]

    This reverts commit 0edabdfe89581669609eaac5f6a8d0ae6fe95e7f.

    I've explained that optional FireWire card for d.2 is also built-in to
    d.2 Pro, however it's wrong. The optional card uses DM1000 ASIC and has
    'Mackie DJ Mixer' in its model name of configuration ROM. On the other
    hand, built-in FireWire card for d.2 Pro and d.4 Pro uses OXFW971 ASIC
    and has 'd.Pro' in its model name according to manuals and user
    experiences. The former card is not the card for d.2 Pro. They are similar
    in appearance but different internally.

    Signed-off-by: Takashi Sakamoto <o-takashi@sakamocchi.jp>
    Link: https://lore.kernel.org/r/20210518084557.102681-2-o-takashi@sakamocchi.jp
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 82e5ee742fdd8874fe996181b87fafe1eb5f1196
Author: Takashi Iwai <tiwai@suse.de>
Date:   Mon May 17 15:15:43 2021 +0200

    ALSA: usx2y: Don't call free_pages_exact() with NULL address

    [ Upstream commit cae0cf651adccee2c3f376e78f30fbd788d0829f ]
```

```
    Unlike some other functions, we can't pass NULL pointer to
    free_pages_exact().  Add a proper NULL check for avoiding possible
    Oops.

    Link: https://lore.kernel.org/r/20210517131545.27252-10-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6c9a4c651572e6806302bb2d1eeebe94e2dcc6d9
Author: Takashi Iwai <tiwai@suse.de>
Date:   Mon May 17 15:15:35 2021 +0200

    ALSA: usx2y: Avoid camelCase

    [ Upstream commit bae3ce4942980d5f7b2b9855f4a2db0c00f9dfbd ]

    For improving readability, convert camelCase fields, variables and
    functions to the plain names with underscore.  Also align the macros
    to be capital letters.

    All done via sed, no functional changes.

    Note that you'll still see many coding style issues even after this
    patch; the fixes will follow.

    Link: https://lore.kernel.org/r/20210517131545.27252-2-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c31e22bf9406c4cc797d4005ffe66f43518e8006
Author: Jonathan Cameron <Jonathan.Cameron@huawei.com>
Date:   Sun May 9 12:33:37 2021 +0100

    iio: magn: bmc150: Balance runtime pm + use pm_runtime_resume_and_get()

    [ Upstream commit 264da512431495e542fcaf56ffe75e7df0e7db74 ]

    probe() error paths after runtime pm is enabled, should disable it.
    remove() should not call pm_runtime_put_noidle() as there is no
    matching get() to have raised the reference count.  This case
    has no affect a the runtime pm core protects against going negative.

    Whilst here use pm_runtime_resume_and_get() to tidy things up a little.
    coccicheck script didn't get this one due to complex code structure so
    found by inspection.

    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Cc: Linus Walleij <linus.walleij@linaro.org>
    Reviewed-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Link: https://lore.kernel.org/r/20210509113354.660190-12-jic23@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7beb9112a3bc58edd2f7566fabaf5001d7eebc14
Author: Jonathan Cameron <Jonathan.Cameron@huawei.com>
Date:   Sun May 9 12:33:27 2021 +0100

    iio: gyro: fxa21002c: Balance runtime pm + use pm_runtime_resume_and_get().

    [ Upstream commit 41120ebbb1eb5e9dec93320e259d5b2c93226073 ]

    In both the probe() error path and remove() pm_runtime_put_noidle()
    is called which will decrement the runtime pm reference count.
    However, there is no matching function to have raised the reference count.
    Not this isn't a fix as the runtime pm core will stop the reference count
```

```
    going negative anyway.

    An alternative would have been to raise the count in these paths, but
    it is not clear why that would be necessary.

    Whilst we are here replace some boilerplate with pm_runtime_resume_and_get()
    Found using coccicheck script under review at:
    https://lore.kernel.org/lkml/20210427141946.2478411-1-Julia.Lawall@inria.fr/

    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Reviewed-by: Rui Miguel Silva <rui.silva@linaro.org>
    Reviewed-by: Mauro Carvalho Chehab <mchehab+huawei@kernel.org>
    Link: https://lore.kernel.org/r/20210509113354.660190-2-jic23@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 78f472928f3a4274dcb7f18cd92ce742fdd6814b
Author: Sean Nyekjaer <sean@geanix.com>
Date:   Fri May 7 12:31:49 2021 +0200

    iio: imu: st_lsm6dsx: correct ODR in header

    [ Upstream commit f7d5c18a8c371c306d73757547c2e0d6cfc764b3 ]

    Fix wrongly stated 13 Hz ODR for accelerometers, the correct ODR is 12.5 Hz

    Signed-off-by: Sean Nyekjaer <sean@geanix.com>
    Acked-by: Lorenzo Bianconi <lorenzo@kernel.org>
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 27a19198ab5833e8434b9234ea3863b6911d8d8f
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Sat May 8 00:07:53 2021 +0200

    partitions: msdos: fix one-byte get_unaligned()

    [ Upstream commit 1b1774998b2dec837a57d729d1a22e5eb2d6d206 ]

    A simplification of get_unaligned() clashes with callers that pass
    in a character pointer, causing a harmless warning like:

    block/partitions/msdos.c: In function 'msdos_partition':
    include/asm-generic/unaligned.h:13:22: warning: 'packed' attribute ignored for field of
type 'u8' {aka 'unsigned char'} [-Wattributes]

    Remove the SYS_IND() macro with the get_unaligned() call
    and just use the ->ind field directly.

    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c029e50e662dce8f0fc32cd7679ed9d30ea5300
Author: Zou Wei <zou_wei@huawei.com>
Date:   Wed May 12 11:54:07 2021 +0800

    ASoC: intel/boards: add missing MODULE_DEVICE_TABLE

    [ Upstream commit a75e5cdf4dd1307bb1541edbb0c008f40896644c ]

    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/1620791647-16024-1-git-send-email-zou_wei@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 09d154990ca82d14aed2b72796f6c8845e2e605d
Author: Tong Zhang <ztong0001@gmail.com>
Date:   Thu May 13 00:07:33 2021 -0400

    misc: alcor_pci: fix null-ptr-deref when there is no PCI bridge

    [ Upstream commit 3ce3e45cc333da707d4d6eb433574b990bcc26f5 ]

    There is an issue with the ASPM(optional) capability checking function.
    A device might be attached to root complex directly, in this case,
    bus->self(bridge) will be NULL, thus priv->parent_pdev is NULL.
    Since alcor_pci_init_check_aspm(priv->parent_pdev) checks the PCI link's
    ASPM capability and populate parent_cap_off, which will be used later by
    alcor_pci_aspm_ctrl() to dynamically turn on/off device, what we can do
    here is to avoid checking the capability if we are on the root complex.
    This will make pdev_cap_off 0 and alcor_pci_aspm_ctrl() will simply
    return when bring called, effectively disable ASPM for the device.

    [    1.246492] BUG: kernel NULL pointer dereference, address: 00000000000000c0
    [    1.248731] RIP: 0010:pci_read_config_byte+0x5/0x40
    [    1.253998] Call Trace:
    [    1.254131]  ? alcor_pci_find_cap_offset.isra.0+0x3a/0x100 [alcor_pci]
    [    1.254476]  alcor_pci_probe+0x169/0x2d5 [alcor_pci]

    Co-developed-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Tong Zhang <ztong0001@gmail.com>
    Link: https://lore.kernel.org/r/20210513040732.1310159-1-ztong0001@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a7268e8a227d5a4f0bd1584f556246b0224ab274
Author: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
Date:   Mon Apr 26 10:06:20 2021 -0700

    misc/libmasm/module: Fix two use after free in ibmasm_init_one

    [ Upstream commit 7272b591c4cb9327c43443f67b8fbae7657dd9ae ]

    In ibmasm_init_one, it calls ibmasm_init_remote_input_dev().
    Inside ibmasm_init_remote_input_dev, mouse_dev and keybd_dev are
    allocated by input_allocate_device(), and assigned to
    sp->remote.mouse_dev and sp->remote.keybd_dev respectively.

    In the err_free_devices error branch of ibmasm_init_one,
    mouse_dev and keybd_dev are freed by input_free_device(), and return
    error. Then the execution runs into error_send_message error branch
    of ibmasm_init_one, where ibmasm_free_remote_input_dev(sp) is called
    to unregister the freed sp->remote.mouse_dev and sp->remote.keybd_dev.

    My patch add a "error_init_remote" label to handle the error of
    ibmasm_init_remote_input_dev(), to avoid the uaf bugs.

    Signed-off-by: Lv Yunlong <lyl2019@mail.ustc.edu.cn>
    Link: https://lore.kernel.org/r/20210426170620.10546-1-lyl2019@mail.ustc.edu.cn
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 221f655abb5693cdb6559f415a04aaf5aee6226c
Author: Jim Quinlan <jim2101024@gmail.com>
Date:   Fri Apr 23 11:32:04 2021 -0700
```

```
    serial: 8250: of: Check for CONFIG_SERIAL_8250_BCM7271

    [ Upstream commit f5b08386dee439c7a9e60ce0a4a4a705f3a60dff ]

    Our SoC's have always had a NS16650A UART core and older SoC's would
    have a compatible string of: 'compatible = ""ns16550a"' and use the
    8250_of driver. Our newer SoC's have added enhancements to the base
    core to add support for DMA and accurate high speed baud rates and use
    this newer 8250_bcm7271 driver. The Device Tree node for our enhanced
    UARTs has a compatible string of: 'compatible = "brcm,bcm7271-uart",
    "ns16550a"''. With both drivers running and the link order setup so
    that the 8250_bcm7217 driver is initialized before the 8250_of driver,
    we should bind the 8250_bcm7271 driver to the enhanced UART, or for
    upstream kernels that don't have the 8250_bcm7271 driver, we bind to
    the 8250_of driver.

    The problem is that when both the 8250_of and 8250_bcm7271 drivers
    were running, occasionally the 8250_of driver would be bound to the
    enhanced UART instead of the 8250_bcm7271 driver. This was happening
    because we use SCMI based clocks which come up late in initialization
    and cause probe DEFER's when the two drivers get their clocks.

    Occasionally the SCMI clock would become ready between the 8250_bcm7271
    probe and the 8250_of probe and the 8250_of driver would be bound. To
    fix this we decided to config only our 8250_bcm7271 driver and added
    "ns16665a0" to the compatible string so the driver would work on our
    older system.

    This commit has of_platform_serial_probe() check specifically for the
    "brcm,bcm7271-uart" and whether its companion driver is enabled. If it
    is the case, and the clock provider is not ready, we want to make sure
    that when the 8250_bcm7271.c driver returns EPROBE_DEFER, we are not
    getting the UART registered via 8250_of.c.

    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Signed-off-by: Jim Quinlan <jim2101024@gmail.com>
    Signed-off-by: Al Cooper <alcooperx@gmail.com>
    Signed-off-by: Florian Fainelli <f.fainelli@gmail.com>
    Link: https://lore.kernel.org/r/20210423183206.3917725-1-f.fainelli@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 464ffcfa6028a6775ca673e57ac08cc4fb6580f1
Author: Michael Walle <michael@walle.cc>
Date:   Wed May 12 16:12:55 2021 +0200

    serial: fsl_lpuart: disable DMA for console and fix sysrq

    [ Upstream commit 8cac2f6eb8548245e6f8fb893fc7f2a714952654 ]

    SYSRQ doesn't work with DMA. This is because there is no error
    indication whether a symbol had a framing error or not. Actually,
    this is not completely correct, there is a bit in the data register
    which is set in this case, but we'd have to read change the DMA access
    to 16 bit and we'd need to post process the data, thus make the DMA
    pointless in the first place.

    Signed-off-by: Michael Walle <michael@walle.cc>
    Link: https://lore.kernel.org/r/20210512141255.18277-10-michael@walle.cc
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c96867d0bba90d69a9961298f3b1820c1086ff69
Author: Sherry Sun <sherry.sun@nxp.com>
Date:   Tue Apr 27 10:12:26 2021 +0800
```

```
    tty: serial: fsl_lpuart: fix the potential risk of division or modulo by zero

    [ Upstream commit fcb10ee27fb91b25b68d7745db9817ecea9f1038 ]

    We should be very careful about the register values that will be used
    for division or modulo operations, althrough the possibility that the
    UARTBAUD register value is zero is very low, but we had better to deal
    with the "bad data" of hardware in advance to avoid division or modulo
    by zero leading to undefined kernel behavior.

    Signed-off-by: Sherry Sun <sherry.sun@nxp.com>
    Link: https://lore.kernel.org/r/20210427021226.27468-1-sherry.sun@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7397dab12fe1e46274f116e7bb74d410ca1123d4
Author: Raymond Tan <raymond.tan@intel.com>
Date:   Wed May 12 16:59:01 2021 +0300

    usb: dwc3: pci: Fix DEFINE for Intel Elkhart Lake

    [ Upstream commit 457d22850b27de3aea336108272d08602c55fdf7 ]

    There's no separate low power (LP) version of Elkhart Lake, thus
    this patch updates the PCI Device ID DEFINE to indicate this.

    Acked-by: Felipe Balbi <balbi@kernel.org>
    Signed-off-by: Raymond Tan <raymond.tan@intel.com>
    Signed-off-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Link: https://lore.kernel.org/r/20210512135901.28495-1-heikki.krogerus@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c4b8190ae66f67866de25689164ff50e5b78d8ef
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Tue May 11 11:00:47 2021 +0800

    soundwire: bus: handle -ENODATA errors in clock stop/start sequences

    [ Upstream commit b50bb8ba369cdc8814e82558117711c12ef3af3d ]

    If a device lost sync and can no longer ACK a command, it may not be
    able to enter a lower-power state but it will still be able to resync
    when the clock restarts. In those cases, we want to continue with the
    clock stop sequence.

    This patch modifies the behavior during clock stop sequences to only
    log errors unrelated to -ENODATA/Command_Ignored. The flow is also
    modified so that loops continue to prepare/deprepare other devices
    even when one seems to have lost sync.

    When resuming the clocks, all issues are logged with a dev_warn(),
    previously only some of them were checked. This is the only part that
    now differs between the clock stop entry and clock stop exit
    sequences: while we don't want to stop the suspend flow, we do want
    information on potential issues while resuming, as they may have
    ripple effects.

    For consistency the log messages are also modified to be unique and
    self-explanatory. Errors in sdw_slave_clk_stop_callback() were
    removed, they are now handled in the caller.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Guennadi Liakhovetski <guennadi.liakhovetski@linux.intel.com>
```

```
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Signed-off-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20210511030048.25622-4-yung-chuan.liao@linux.intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 740b019c2eb80642e3a2609bf7450f0b267cf4c4
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Tue May 11 11:00:45 2021 +0800

    soundwire: bus: only use CLOCK_STOP_MODE0 and fix confusions

    [ Upstream commit 345e9f5ca798600e44c0843646621f2804eb99f4 ]

    Existing devices and implementations only support the required
    CLOCK_STOP_MODE0. All the code related to CLOCK_STOP_MODE1 has not
    been tested and is highly questionable, with a clear confusion between
    CLOCK_STOP_MODE1 and the simple clock stop state machine.

    This patch removes all usages of CLOCK_STOP_MODE1 - which has no
    impact on any solution - and fixes the use of the simple clock stop
    state machine. The resulting code should be a lot more symmetrical and
    easier to maintain.

    Note that CLOCK_STOP_MODE1 is not supported in the SoundWire Device
    Class specification so it's rather unlikely that we need to re-add
    this mode later.

    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Guennadi Liakhovetski <guennadi.liakhovetski@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Signed-off-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Link: https://lore.kernel.org/r/20210511030048.25622-2-yung-chuan.liao@linux.intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee188d0f8dd59c5db9f4ad761c97dd30e9d6361d
Author: Paul E. McKenney <paulmck@kernel.org>
Date:   Mon Apr 5 09:51:05 2021 -0700

    rcu: Reject RCU_LOCKDEP_WARN() false positives

    [ Upstream commit 3066820034b5dd4e89bd74a7739c51c2d6f5e554 ]

    If another lockdep report runs concurrently with an RCU lockdep report
    from RCU_LOCKDEP_WARN(), the following sequence of events can occur:

    1.      debug_lockdep_rcu_enabled() sees that lockdep is enabled
            when called from (say) synchronize_rcu().

    2.      Lockdep is disabled by a concurrent lockdep report.

    3.      debug_lockdep_rcu_enabled() evaluates its lockdep-expression
            argument, for example, lock_is_held(&rcu_bh_lock_map).

    4.      Because lockdep is now disabled, lock_is_held() plays it safe and
            returns the constant 1.

    5.      But in this case, the constant 1 is not safe, because invoking
            synchronize_rcu() under rcu_read_lock_bh() is disallowed.

    6.      debug_lockdep_rcu_enabled() wrongly invokes lockdep_rcu_suspicious(),
            resulting in a false-positive splat.

    This commit therefore changes RCU_LOCKDEP_WARN() to check
```

```
    debug_lockdep_rcu_enabled() after checking the lockdep expression,
    so that any "safe" returns from lock_is_held() are rejected by
    debug_lockdep_rcu_enabled().  This requires memory ordering, which is
    supplied by READ_ONCE(debug_locks).  The resulting volatile accesses
    prevent the compiler from reordering and the fact that only one variable
    is being accessed prevents the underlying hardware from reordering.
    The combination works for IA64, which can reorder reads to the same
    location, but this is defeated by the volatile accesses, which compile
    to load instructions that provide ordering.

    Reported-by: syzbot+dde0cc33951735441301@syzkaller.appspotmail.com
    Reported-by: Matthew Wilcox <willy@infradead.org>
    Reported-by: syzbot+88e4f02896967fe1ab0d@syzkaller.appspotmail.com
    Reported-by: Thomas Gleixner <tglx@linutronix.de>
    Suggested-by: Boqun Feng <boqun.feng@gmail.com>
    Reviewed-by: Boqun Feng <boqun.feng@gmail.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa59856bcfd20107b4855c28e863b27be26f6816
Author: Frederic Weisbecker <frederic@kernel.org>
Date:   Sat Apr 17 15:16:49 2021 +0200

    srcu: Fix broken node geometry after early ssp init

    [ Upstream commit b5befe842e6612cf894cf4a199924ee872d8b7d8 ]

    An srcu_struct structure that is initialized before rcu_init_geometry()
    will have its srcu_node hierarchy based on CONFIG_NR_CPUS.  Once
    rcu_init_geometry() is called, this hierarchy is compressed as needed
    for the actual maximum number of CPUs for this system.

    Later on, that srcu_struct structure is confused, sometimes referring
    to its initial CONFIG_NR_CPUS-based hierarchy, and sometimes instead
    to the new num_possible_cpus() hierarchy.  For example, each of its
    ->mynode fields continues to reference the original leaf rcu_node
    structures, some of which might no longer exist.  On the other hand,
    srcu_for_each_node_breadth_first() traverses to the new node hierarchy.

    There are at least two bad possible outcomes to this:

    1) a) A callback enqueued early on an srcu_data structure (call it
          *sdp) is recorded pending on sdp->mynode->srcu_data_have_cbs in
          srcu_funnel_gp_start() with sdp->mynode pointing to a deep leaf
          (say 3 levels).

       b) The grace period ends after rcu_init_geometry() shrinks the
          nodes level to a single one.  srcu_gp_end() walks through the new
          srcu_node hierarchy without ever reaching the old leaves so the
          callback is never executed.

       This is easily reproduced on an 8 CPUs machine with CONFIG_NR_CPUS >= 32
       and "rcupdate.rcu_self_test=1". The srcu_barrier() after early tests
       verification never completes and the boot hangs:

            [ 5413.141029] INFO: task swapper/0:1 blocked for more than 4915 seconds.
            [ 5413.147564]       Not tainted 5.12.0-rc4+ #28
            [ 5413.151927] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this
message.
            [ 5413.159753] task:swapper/0       state:D stack:    0 pid:    1 ppid:     0
flags:0x00004000
            [ 5413.168099] Call Trace:
            [ 5413.170555]  __schedule+0x36c/0x930
            [ 5413.174057]  ? wait_for_completion+0x88/0x110
            [ 5413.178423]  schedule+0x46/0xf0
```

```
              [ 5413.181575]  schedule_timeout+0x284/0x380
              [ 5413.185591]  ? wait_for_completion+0x88/0x110
              [ 5413.189957]  ? mark_held_locks+0x61/0x80
              [ 5413.193882]  ? mark_held_locks+0x61/0x80
              [ 5413.197809]  ? _raw_spin_unlock_irq+0x24/0x50
              [ 5413.202173]  ? wait_for_completion+0x88/0x110
              [ 5413.206535]  wait_for_completion+0xb4/0x110
              [ 5413.210724]  ? srcu_torture_stats_print+0x110/0x110
              [ 5413.215610]  srcu_barrier+0x187/0x200
              [ 5413.219277]  ? rcu_tasks_verify_self_tests+0x50/0x50
              [ 5413.224244]  ? rdinit_setup+0x2b/0x2b
              [ 5413.227907]  rcu_verify_early_boot_tests+0x2d/0x40
              [ 5413.232700]  do_one_initcall+0x63/0x310
              [ 5413.236541]  ? rdinit_setup+0x2b/0x2b
              [ 5413.240207]  ? rcu_read_lock_sched_held+0x52/0x80
              [ 5413.244912]  kernel_init_freeable+0x253/0x28f
              [ 5413.249273]  ? rest_init+0x250/0x250
              [ 5413.252846]  kernel_init+0xa/0x110
              [ 5413.256257]  ret_from_fork+0x22/0x30

    2) An srcu_struct structure that is initialized before rcu_init_geometry()
       and used afterward will always have stale rdp->mynode references,
       resulting in callbacks to be missed in srcu_gp_end(), just like in
       the previous scenario.

    This commit therefore causes init_srcu_struct_nodes to initialize the
    geometry, if needed.  This ensures that the srcu_node hierarchy is
    properly built and distributed from the get-go.

    Suggested-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Frederic Weisbecker <frederic@kernel.org>
    Cc: Boqun Feng <boqun.feng@gmail.com>
    Cc: Lai Jiangshan <jiangshanlai@gmail.com>
    Cc: Neeraj Upadhyay <neeraju@codeaurora.org>
    Cc: Josh Triplett <josh@joshtriplett.org>
    Cc: Joel Fernandes <joel@joelfernandes.org>
    Cc: Uladzislau Rezki <urezki@gmail.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5318059dea5e1c95671c712cf020fa1ef7008526
Author: ching Huang <ching2048@areca.com.tw>
Date:   Fri Apr 16 11:44:57 2021 +0800

    scsi: arcmsr: Fix the wrong CDB payload report to IOP

    [ Upstream commit 5b8644968d2ca85abb785e83efec36934974b0c2 ]

    This patch fixes the wrong CDB payload report to IOP.

    Link:
https://lore.kernel.org/r/d2c97df3c817595c6faf582839316209022f70da.camel@areca.com.tw
    Signed-off-by: ching Huang <ching2048@areca.com.tw>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 38a6e9b501b24ffb262d383a3522ad304b2b8469
Author: Robin Gong <yibin.gong@nxp.com>
Date:   Mon Apr 26 16:59:09 2021 +0800

    dmaengine: fsl-qdma: check dma_set_mask return value

    [ Upstream commit f0c07993af0acf5545d5c1445798846565f4f147 ]

    For fix below warning reported by static code analysis tool like Coverity
```

```
      from Synopsys:

      Signed-off-by: Robin Gong <yibin.gong@nxp.com>
      Addresses-Coverity-ID: 12285639 ("Unchecked return value")
      Link: https://lore.kernel.org/r/1619427549-20498-1-git-send-email-yibin.gong@nxp.com
      Signed-off-by: Vinod Koul <vkoul@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 7bd674051bfec6544ca5693d5d58119aa510de31
  Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
  Date:   Wed May 5 11:36:56 2021 -0500

      ASoC: Intel: sof_sdw: add mutual exclusion between PCH DMIC and RT715

      [ Upstream commit 35564e2bf94611c3eb51d35362addb3cb394ad54 ]

      When external RT714/715 devices are used for capture, we don't want
      the PCH DMICs to be used.

      Any information provided by the SOF platform driver or DMI quirks will
      be overridden.

      Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
      Reviewed-by: Kai Vehmanen <kai.vehmanen@linux.intel.com>
      Reviewed-by: Libin Yang <libin.yang@intel.com>
      Link: https://lore.kernel.org/r/20210505163705.305616-5-pierre-
  louis.bossart@linux.intel.com
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c57e7b2f36ae0d8fb88b5e71fb33a29987939b78
  Author: Yang Yingliang <yangyingliang@huawei.com>
  Date:   Tue Apr 6 21:11:59 2021 +0800

      leds: tlc591xx: fix return value check in tlc591xx_probe()

      [ Upstream commit ee522bcf026ec82ada793979c3a906274430595a ]

      After device_get_match_data(), tlc591xx is not checked, add
      check for it and also check np after dev_of_node.

      Reported-by: Hulk Robot <hulkci@huawei.com>
      Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
      Signed-off-by: Pavel Machek <pavel@ucw.cz>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 28f0adaae3e8167f9355c7c0b8020de7e8d981d9
  Author: Nikolay Aleksandrov <nikolay@nvidia.com>
  Date:   Sun Jul 11 12:56:29 2021 +0300

      net: bridge: multicast: fix MRD advertisement router port marking race

      commit 000b7287b67555fee39d39fff75229dedde0dcbf upstream.

      When an MRD advertisement is received on a bridge port with multicast
      snooping enabled, we mark it as a router port automatically, that
      includes adding that port to the router port list. The multicast lock
      protects that list, but it is not acquired in the MRD advertisement case
      leading to a race condition, we need to take it to fix the race.

      Cc: stable@vger.kernel.org
      Cc: linus.luessing@c0d3.blue
      Fixes: 4b3087c7e37f ("bridge: Snoop Multicast Router Advertisements")
      Signed-off-by: Nikolay Aleksandrov <nikolay@nvidia.com>
      Signed-off-by: David S. Miller <davem@davemloft.net>
```

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

commit dc924e41187097a2e8baf8097377137d23bac1c7
Author: Nikolay Aleksandrov <nikolay@nvidia.com>
Date:   Sun Jul 11 12:56:28 2021 +0300

    net: bridge: multicast: fix PIM hello router port marking race

    commit 04bef83a3358946bfc98a5ecebd1b0003d83d882 upstream.

    When a PIM hello packet is received on a bridge port with multicast
    snooping enabled, we mark it as a router port automatically, that
    includes adding that port the router port list. The multicast lock
    protects that list, but it is not acquired in the PIM message case
    leading to a race condition, we need to take it to fix the race.

    Cc: stable@vger.kernel.org
    Fixes: 91b02d3d133b ("bridge: mcast: add router port on PIM hello message")
    Signed-off-by: Nikolay Aleksandrov <nikolay@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b937eb91b35316206e8ec7ace44852f4b28e2142
Author: José Roberto de Souza <jose.souza@intel.com>
Date:   Wed Jun 16 12:44:15 2021 -0700

    drm/dp_mst: Add missing drm parameters to recently added call to drm_dbg_kms()

    commit 24ff3dc18b99c4b912ab1746e803ddb3be5ced4c upstream.

    Commit 3769e4c0af5b ("drm/dp_mst: Avoid to mess up payload table by
    ports in stale topology") added to calls to drm_dbg_kms() but it
    missed the first parameter, the drm device breaking the build.

    Fixes: 3769e4c0af5b ("drm/dp_mst: Avoid to mess up payload table by ports in stale
topology")
    Cc: Wayne Lin <Wayne.Lin@amd.com>
    Cc: Lyude Paul <lyude@redhat.com>
    Cc: dri-devel@lists.freedesktop.org
    Cc: stable@vger.kernel.org
    Signed-off-by: José Roberto de Souza <jose.souza@intel.com>
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Signed-off-by: Lyude Paul <lyude@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210616194415.36926-1-
jose.souza@intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d3891139901bb4e6f66defc10b2a8c0c4d5b2e9f
Author: Wayne Lin <Wayne.Lin@amd.com>
Date:   Wed Jun 16 11:55:01 2021 +0800

    drm/dp_mst: Avoid to mess up payload table by ports in stale topology

    commit 3769e4c0af5b82c8ea21d037013cb9564dfaa51f upstream.

    [Why]
    After unplug/hotplug hub from the system, userspace might start to
    clear stale payloads gradually. If we call drm_dp_mst_deallocate_vcpi()
    to release stale VCPI of those ports which are not relating to current
    topology, we have chane to wrongly clear active payload table entry for
    current topology.

    E.g.
    We have allocated VCPI 1 in current payload table and we call
    drm_dp_mst_deallocate_vcpi() to clean VCPI 1 in stale topology. In
```

```
    drm_dp_mst_deallocate_vcpi(), it will call drm_dp_mst_put_payload_id()
    tp put VCPI 1 and which means ID 1 is available again. Thereafter, if we
    want to allocate a new payload stream, it will find ID 1 is available by
    drm_dp_mst_assign_payload_id(). However, ID 1 is being used

    [How]
    Check target sink is relating to current topology or not before doing
    any payload table update.
    Searching upward to find the target sink's relevant root branch device.
    If the found root branch device is not the same root of current
    topology, don't update payload table.

    Changes since v1:
    * Change debug macro to use drm_dbg_kms() instead
    * Amend the commit message to add Cc tag.

    Signed-off-by: Wayne Lin <Wayne.Lin@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Lyude Paul <lyude@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210616035501.3776-3-
Wayne.Lin@amd.com
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d01943939f45d2beac7a561489eefe45a08a94a6
Author: Wayne Lin <Wayne.Lin@amd.com>
Date:   Wed Jun 16 11:55:00 2021 +0800

    drm/dp_mst: Do not set proposed vcpi directly

    commit 35d3e8cb35e75450f87f87e3d314e2d418b6954b upstream.

    [Why]
    When we receive CSN message to notify one port is disconnected, we will
    implicitly set its corresponding num_slots to 0. Later on, we will
    eventually call drm_dp_update_payload_part1() to arrange down streams.

    In drm_dp_update_payload_part1(), we iterate over all proposed_vcpis[]
    to do the update. Not specific to a target sink only. For example, if we
    light up 2 monitors, Monitor_A and Monitor_B, and then we unplug
    Monitor_B. Later on, when we call drm_dp_update_payload_part1() to try
    to update payload for Monitor_A, we'll also implicitly clean payload for
    Monitor_B at the same time. And finally, when we try to call
    drm_dp_update_payload_part1() to clean payload for Monitor_B, we will do
    nothing at this time since payload for Monitor_B has been cleaned up
    previously.

    For StarTech 1to3 DP hub, it seems like if we didn't update DPCD payload
    ID table then polling for "ACT Handled"(BIT_1 of DPCD 002C0h) will fail
    and this polling will last for 3 seconds.

    Therefore, guess the best way is we don't set the proposed_vcpi[]
    diretly. Let user of these herlper functions to set the proposed_vcpi
    directly.

    [How]
    1. Revert commit 7617e9621bf2 ("drm/dp_mst: clear time slots for ports
    invalid")
    2. Tackle the issue in previous commit by skipping those trasient
    proposed VCPIs. These stale VCPIs shoulde be explicitly cleared by
    user later on.

    Changes since v1:
    * Change debug macro to use drm_dbg_kms() instead
    * Amend the commit message to add Fixed & Cc tags
```

```
    Signed-off-by: Wayne Lin <Wayne.Lin@amd.com>
    Fixes: 7617e9621bf2 ("drm/dp_mst: clear time slots for ports invalid")
    Cc: Lyude Paul <lyude@redhat.com>
    Cc: Wayne Lin <Wayne.Lin@amd.com>
    Cc: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    Cc: Maxime Ripard <mripard@kernel.org>
    Cc: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: dri-devel@lists.freedesktop.org
    Cc: <stable@vger.kernel.org> # v5.5+
    Signed-off-by: Lyude Paul <lyude@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210616035501.3776-2-
Wayne.Lin@amd.com
    Reviewed-by: Lyude Paul <lyude@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c30f0b5c6fd1e023d16033a9c90f8ec166deaef3
Author: Filipe Manana <fdmanana@suse.com>
Date:   Wed Jul 7 12:23:45 2021 +0100

    btrfs: zoned: fix wrong mutex unlock on failure to allocate log root tree

    commit ea32af47f00a046a1f953370514d6d946efe0152 upstream.

    When syncing the log, if we fail to allocate the root node for the log
    root tree:

    1) We are unlocking fs_info->tree_log_mutex, but at this point we have
       not yet locked this mutex;

    2) We have locked fs_info->tree_root->log_mutex, but we end up not
       unlocking it;

    So fix this by unlocking fs_info->tree_root->log_mutex instead of
    fs_info->tree_log_mutex.

    Fixes: e75f9fd194090e ("btrfs: zoned: move log tree node allocation out of
log_root_tree->log_mutex")
    CC: stable@vger.kernel.org # 5.13+
    Reviewed-by: Nikolay Borisov <nborisov@suse.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 02fac1508a726d3cec4ef71b66e0a7ce65521b39
Author: Johannes Thumshirn <johannes.thumshirn@wdc.com>
Date:   Tue Jul 6 01:32:38 2021 +0900

    btrfs: don't block if we can't acquire the reclaim lock

    commit 9cc0b837e14ae913581ec1ea6e979a738f71b0fd upstream.

    If we can't acquire the reclaim_bgs_lock on block group reclaim, we
    block until it is free. This can potentially stall for a long time.

    While reclaim of block groups is necessary for a good user experience on
    a zoned file system, there still is no need to block as it is best
    effort only, just like when we're deleting unused block groups.

    CC: stable@vger.kernel.org # 5.13
    Signed-off-by: Johannes Thumshirn <johannes.thumshirn@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 d6479f089af166ce8696a14fed3d73ea378d0173
Author: Filipe Manana <fdmanana@suse.com>
Date:   Tue Jun 29 14:43:06 2021 +0100

    btrfs: rework chunk allocation to avoid exhaustion of the system chunk array

    commit 79bd37120b149532af5b21953643ed74af69654f upstream.

    Commit eafa4fd0ad0607 ("btrfs: fix exhaustion of the system chunk array
    due to concurrent allocations") fixed a problem that resulted in
    exhausting the system chunk array in the superblock when there are many
    tasks allocating chunks in parallel. Basically too many tasks enter the
    first phase of chunk allocation without previous tasks having finished
    their second phase of allocation, resulting in too many system chunks
    being allocated. That was originally observed when running the fallocate
    tests of stress-ng on a PowerPC machine, using a node size of 64K.

    However that commit also introduced a deadlock where a task in phase 1 of
    the chunk allocation waited for another task that had allocated a system
    chunk to finish its phase 2, but that other task was waiting on an extent
    buffer lock held by the first task, therefore resulting in both tasks not
    making any progress. That change was later reverted by a patch with the
    subject "btrfs: fix deadlock with concurrent chunk allocations involving
    system chunks", since there is no simple and short solution to address it
    and the deadlock is relatively easy to trigger on zoned filesystems, while
    the system chunk array exhaustion is not so common.

    This change reworks the chunk allocation to avoid the system chunk array
    exhaustion. It accomplishes that by making the first phase of chunk
    allocation do the updates of the device items in the chunk btree and the
    insertion of the new chunk item in the chunk btree. This is done while
    under the protection of the chunk mutex (fs_info->chunk_mutex), in the
    same critical section that checks for available system space, allocates
    a new system chunk if needed and reserves system chunk space. This way
    we do not have chunk space reserved until the second phase completes.

    The same logic is applied to chunk removal as well, since it keeps
    reserved system space long after it is done updating the chunk btree.

    For direct allocation of system chunks, the previous behaviour remains,
    because otherwise we would deadlock on extent buffers of the chunk btree.
    Changes to the chunk btree are by large done by chunk allocation and chunk
    removal, which first reserve chunk system space and then later do changes
    to the chunk btree. The other remaining cases are uncommon and correspond
    to adding a device, removing a device and resizing a device. All these
    other cases do not pre-reserve system space, they modify the chunk btree
    right away, so they don't hold reserved space for a long period like chunk
    allocation and chunk removal do.

    The diff of this change is huge, but more than half of it is just addition
    of comments describing both how things work regarding chunk allocation and
    removal, including both the new behavior and the parts of the old behavior
    that did not change.

    CC: stable@vger.kernel.org # 5.12+
    Tested-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Tested-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Tested-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 789b24d9950d3e67b227f81b3fab912a8fb257af
Author: Filipe Manana <fdmanana@suse.com>
```

Date: Tue Jun 29 14:43:05 2021 +0100

    btrfs: fix deadlock with concurrent chunk allocations involving system chunks

    commit 1cb3db1cf383a3c7dbda1aa0ce748b0958759947 upstream.

    When a task attempting to allocate a new chunk verifies that there is not
    currently enough free space in the system space_info and there is another
    task that allocated a new system chunk but it did not finish yet the
    creation of the respective block group, it waits for that other task to
    finish creating the block group. This is to avoid exhaustion of the system
    chunk array in the superblock, which is limited, when we have a thundering
    herd of tasks allocating new chunks. This problem was described and fixed
    by commit eafa4fd0ad0607 ("btrfs: fix exhaustion of the system chunk array
    due to concurrent allocations").

    However there are two very similar scenarios where this can lead to a
    deadlock:

    1) Task B allocated a new system chunk and task A is waiting on task B
       to finish creation of the respective system block group. However before
       task B ends its transaction handle and finishes the creation of the
       system block group, it attempts to allocate another chunk (like a data
       chunk for an fallocate operation for a very large range). Task B will
       be unable to progress and allocate the new chunk, because task A set
       space_info->chunk_alloc to 1 and therefore it loops at
       btrfs_chunk_alloc() waiting for task A to finish its chunk allocation
       and set space_info->chunk_alloc to 0, but task A is waiting on task B
       to finish creation of the new system block group, therefore resulting
       in a deadlock;

    2) Task B allocated a new system chunk and task A is waiting on task B to
       finish creation of the respective system block group. By the time that
       task B enter the final phase of block group allocation, which happens
       at btrfs_create_pending_block_groups(), when it modifies the extent
       tree, the device tree or the chunk tree to insert the items for some
       new block group, it needs to allocate a new chunk, so it ends up at
       btrfs_chunk_alloc() and keeps looping there because task A has set
       space_info->chunk_alloc to 1, but task A is waiting for task B to
       finish creation of the new system block group and release the reserved
       system space, therefore resulting in a deadlock.

    In short, the problem is if a task B needs to allocate a new chunk after
    it previously allocated a new system chunk and if another task A is
    currently waiting for task B to complete the allocation of the new system
    chunk.

    Unfortunately this deadlock scenario introduced by the previous fix for
    the system chunk array exhaustion problem does not have a simple and short
    fix, and requires a big change to rework the chunk allocation code so that
    chunk btree updates are all made in the first phase of chunk allocation.
    And since this deadlock regression is being frequently hit on zoned
    filesystems and the system chunk array exhaustion problem is triggered
    in more extreme cases (originally observed on PowerPC with a node size
    of 64K when running the fallocate tests from stress-ng), revert the
    changes from that commit. The next patch in the series, with a subject
    of "btrfs: rework chunk allocation to avoid exhaustion of the system
    chunk array" does the necessary changes to fix the system chunk array
    exhaustion problem.

    Reported-by: Naohiro Aota <naohiro.aota@wdc.com>
    Link: https://lore.kernel.org/linux-btrfs/20210621015922.ewgbffxuawia7liz@naota-xeon/
    Fixes: eafa4fd0ad0607 ("btrfs: fix exhaustion of the system chunk array due to
concurrent allocations")
    CC: stable@vger.kernel.org # 5.12+

```
    Tested-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Tested-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: Filipe Manana <fdmanana@suse.com>
    Tested-by: David Sterba <dsterba@suse.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9ffb4cd20998ad5038aa4a4c17e894cc2d37def3
Author: David Sterba <dsterba@suse.com>
Date:   Wed Jun 23 17:54:54 2021 +0200

    btrfs: zoned: fix types for u64 division in btrfs_reclaim_bgs_work

    commit 54afaae34ee49e98c1c902b444b42832551d090c upstream.

    The types in calculation of the used percentage in the reclaiming
    messages are both u64, though bg->length is either 1GiB (non-zoned) or
    the zone size in the zoned mode. The upper limit on zone size is 8GiB so
    this could theoretically overflow in the future, right now the values
    fit.

    Fixes: 18bb8bbf13c1 ("btrfs: zoned: automatically reclaim zones")
    CC: stable@vger.kernel.org # 5.13
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 27c5ecbb691fb89ee5753b109be432353bddfcaf
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Mon Jun 28 17:57:28 2021 +0900

    btrfs: properly split extent_map for REQ_OP_ZONE_APPEND

    commit abb99cfdaf0759f8a619e5fecf52ccccdf310c8c upstream.

    Damien reported a test failure with btrfs/209. The test itself ran fine,
    but the fsck ran afterwards reported a corrupted filesystem.

    The filesystem corruption happens because we're splitting an extent and
    then writing the extent twice. We have to split the extent though, because
    we're creating too large extents for a REQ_OP_ZONE_APPEND operation.

    When dumping the extent tree, we can see two EXTENT_ITEMs at the same
    start address but different lengths.

    $ btrfs inspect dump-tree /dev/nullb1 -t extent
    ...
       item 19 key (269484032 EXTENT_ITEM 126976) itemoff 15470 itemsize 53
               refs 1 gen 7 flags DATA
               extent data backref root FS_TREE objectid 257 offset 786432 count 1
       item 20 key (269484032 EXTENT_ITEM 262144) itemoff 15417 itemsize 53
               refs 1 gen 7 flags DATA
               extent data backref root FS_TREE objectid 257 offset 786432 count 1

    The duplicated EXTENT_ITEMs originally come from wrongly split extent_map in
    extract_ordered_extent(). Since extract_ordered_extent() uses
    create_io_em() to split an existing extent_map, we will have
    split->orig_start != split->start. Then, it will be logged with non-zero
    "extent data offset". Finally, the logged entries are replayed into
    a duplicated EXTENT_ITEM.

    Introduce and use proper splitting function for extent_map. The function is
    intended to be simple and specific usage for extract_ordered_extent() e.g.
    not supporting compression case (we do not allow splitting compressed
    extent_map anyway).
```

```
    There was a question raised by Qu, in summary why we want to split the
    extent map (and not the bio):

    The problem is not the limit on the zone end, which as you mention is
    the same as the block group end. The problem is that data write use zone
    append (ZA) operations. ZA BIOs cannot be split so a large extent may
    need to be processed with multiple ZA BIOs, While that is also true for
    regular writes, the major difference is that ZA are "nameless" write
    operation giving back the written sectors on completion. And ZA
    operations may be reordered by the block layer (not intentionally
    though). Combine both of these characteristics and you can see that the
    data for a large extent may end up being shuffled when written resulting
    in data corruption and the impossibility to map the extent to some start
    sector.

    To avoid this problem, zoned btrfs uses the principle "one data extent
    == one ZA BIO". So large extents need to be split. This is unfortunate,
    but we can revisit this later and optimize, e.g. merge back together the
    fragments of an extent once written if they actually were written
    sequentially in the zone.

    Reported-by: Damien Le Moal <damien.lemoal@wdc.com>
    Fixes: d22002fd37bd ("btrfs: zoned: split ordered extent when bio is sent")
    CC: stable@vger.kernel.org # 5.12+
    CC: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 08e1e3b6ddd89cb3283f1da071258e781282a068
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Sat Jul 10 02:45:59 2021 +0100

    io_uring: use right task for exiting checks

    commit 9c6882608bce249a8918744ecdb65748534e3f17 upstream.

    When we use delayed_work for fallback execution of requests, current
    will be not of the submitter task, and so checks in io_req_task_submit()
    may not behave as expected. Currently, it leaves inline completions not
    flushed, so making io_ring_exit_work() to hang. Use the submitter task
    for all those checks.

    Cc: stable@vger.kernel.org
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/cb413c715bed0bc9c98b169059ea9c8a2c770715.1625881431.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6ef21f34aa92d05f4b706fd954451fffc658735f
Author: Robin Murphy <robin.murphy@arm.com>
Date:   Mon Jul 12 15:27:46 2021 +0100

    arm64: Avoid premature usercopy failure

    commit 295cf156231ca3f9e3a66bde7fab5e09c41835e0 upstream.

    Al reminds us that the usercopy API must only return complete failure
    if absolutely nothing could be copied. Currently, if userspace does
    something silly like giving us an unaligned pointer to Device memory,
    or a size which overruns MTE tag bounds, we may fail to honour that
    requirement when faulting on a multi-byte access even though a smaller
```

```
    access could have succeeded.

    Add a mitigation to the fixup routines to fall back to a single-byte
    copy if we faulted on a larger access before anything has been written
    to the destination, to guarantee making *some* forward progress. We
    needn't be too concerned about the overall performance since this should
    only occur when callers are doing something a bit dodgy in the first
    place. Particularly broken userspace might still be able to trick
    generic_perform_write() into an infinite loop by targeting write() at
    an mmap() of some read-only device register where the fault-in load
    succeeds but any store synchronously aborts such that copy_to_user() is
    genuinely unable to make progress, but, well, don't do that...

    CC: stable@vger.kernel.org
    Reported-by: Chen Huang <chenhuang5@huawei.com>
    Suggested-by: Al Viro <viro@zeniv.linux.org.uk>
    Reviewed-by: Catalin Marinas <catalin.marinas@arm.com>
    Signed-off-by: Robin Murphy <robin.murphy@arm.com>
    Link:
https://lore.kernel.org/r/dc03d5c675731a1f24a62417dba5429ad744234e.1626098433.git.robin.murp
hy@arm.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 70fa723908a42fdc84d5a93dbac05aedd015b345
Author: Joao Martins <joao.m.martins@oracle.com>
Date:   Wed Jul 14 21:27:11 2021 -0700

    mm/hugetlb: fix refs calculation from unaligned @vaddr

    commit d08af0a59684e18a51aa4bfd24c658994ea3fc5b upstream.

    Commit 82e5d378b0e47 ("mm/hugetlb: refactor subpage recording")
    refactored the count of subpages but missed an edge case when @vaddr is
    not aligned to PAGE_SIZE e.g.  when close to vma->vm_end.  It would then
    errousnly set @refs to 0 and record_subpages_vmas() wouldn't set the
    @pages array element to its value, consequently causing the reported
    null-deref by syzbot.

    Fix it by aligning down @vaddr by PAGE_SIZE in @refs calculation.

    Link: https://lkml.kernel.org/r/20210713152440.28650-1-joao.m.martins@oracle.com
    Fixes: 82e5d378b0e47 ("mm/hugetlb: refactor subpage recording")
    Reported-by: syzbot+a3fcd59df1b372066f5a@syzkaller.appspotmail.com
    Signed-off-by: Joao Martins <joao.m.martins@oracle.com>
    Reviewed-by: Mike Kravetz <mike.kravetz@oracle.com>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 65a5977bd05c3dc6753b0032fa557a9539d99a9b
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Thu Jul 15 11:55:31 2021 -0700

    EDAC/igen6: fix core dependency AGAIN

    commit a1c9ca5f65c9acfd7c02474b9d5cacbd7ea288df upstream.

    My previous patch had a typo/thinko which prevents this driver
    from being enabled: change X64_64 to X86_64.

    Fixes: 0a9ece9ba154 ("EDAC/igen6: fix core dependency")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
```

```
      Cc: Borislav Petkov <bp@alien8.de>
      Cc: Mauro Carvalho Chehab <mchehab@kernel.org>
      Cc: linux-edac@vger.kernel.org
      Cc: bowsingbetee <bowsingbetee@protonmail.com>
      Cc: stable@vger.kernel.org
      Signed-off-by: Tony Luck <tony.luck@intel.com>
      Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit d6e76469157d8f240e5dec6f8411aa8d306b1126
  Author: Zhen Lei <thunder.leizhen@huawei.com>
  Date:   Mon Jul 12 16:55:44 2021 +0800

      fbmem: Do not delete the mode that is still in use

      commit 0af778269a522c988ef0b4188556aba97fb420cc upstream.

      The execution of fb_delete_videomode() is not based on the result of the
      previous fbcon_mode_deleted(). As a result, the mode is directly deleted,
      regardless of whether it is still in use, which may cause UAF.

      ==================================================================
      BUG: KASAN: use-after-free in fb_mode_is_equal+0x36e/0x5e0 \
      drivers/video/fbdev/core/modedb.c:924
      Read of size 4 at addr ffff88807e0ddb1c by task syz-executor.0/18962

      CPU: 2 PID: 18962 Comm: syz-executor.0 Not tainted 5.10.45-rc1+ #3
      Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS ...
      Call Trace:
       __dump_stack lib/dump_stack.c:77 [inline]
       dump_stack+0x137/0x1be lib/dump_stack.c:118
       print_address_description+0x6c/0x640 mm/kasan/report.c:385
       __kasan_report mm/kasan/report.c:545 [inline]
       kasan_report+0x13d/0x1e0 mm/kasan/report.c:562
       fb_mode_is_equal+0x36e/0x5e0 drivers/video/fbdev/core/modedb.c:924
       fbcon_mode_deleted+0x16a/0x220 drivers/video/fbdev/core/fbcon.c:2746
       fb_set_var+0x1e1/0xdb0 drivers/video/fbdev/core/fbmem.c:975
       do_fb_ioctl+0x4d9/0x6e0 drivers/video/fbdev/core/fbmem.c:1108
       vfs_ioctl fs/ioctl.c:48 [inline]
       __do_sys_ioctl fs/ioctl.c:753 [inline]
       __se_sys_ioctl+0xfb/0x170 fs/ioctl.c:739
       do_syscall_64+0x2d/0x70 arch/x86/entry/common.c:46
       entry_SYSCALL_64_after_hwframe+0x44/0xa9

      Freed by task 18960:
       kasan_save_stack mm/kasan/common.c:48 [inline]
       kasan_set_track+0x3d/0x70 mm/kasan/common.c:56
       kasan_set_free_info+0x17/0x30 mm/kasan/generic.c:355
       __kasan_slab_free+0x108/0x140 mm/kasan/common.c:422
       slab_free_hook mm/slub.c:1541 [inline]
       slab_free_freelist_hook+0xd6/0x1a0 mm/slub.c:1574
       slab_free mm/slub.c:3139 [inline]
       kfree+0xca/0x3d0 mm/slub.c:4121
       fb_delete_videomode+0x56a/0x820 drivers/video/fbdev/core/modedb.c:1104
       fb_set_var+0x1f3/0xdb0 drivers/video/fbdev/core/fbmem.c:978
       do_fb_ioctl+0x4d9/0x6e0 drivers/video/fbdev/core/fbmem.c:1108
       vfs_ioctl fs/ioctl.c:48 [inline]
       __do_sys_ioctl fs/ioctl.c:753 [inline]
       __se_sys_ioctl+0xfb/0x170 fs/ioctl.c:739
       do_syscall_64+0x2d/0x70 arch/x86/entry/common.c:46
       entry_SYSCALL_64_after_hwframe+0x44/0xa9

      Fixes: 13ff178ccd6d ("fbcon: Call fbcon_mode_deleted/new_modelist directly")
      Signed-off-by: Zhen Lei <thunder.leizhen@huawei.com>
      Cc: <stable@vger.kernel.org> # v5.3+
```

```
    Signed-off-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210712085544.2828-1-
thunder.leizhen@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a41573667b39152176f6b08d10b4deb171e541c4
Author: Christian Brauner <christian.brauner@ubuntu.com>
Date:   Wed Jul 14 15:47:49 2021 +0200

    cgroup: verify that source is a string

    commit 3b0462726e7ef281c35a7a4ae33e93ee2bc9975b upstream.

    The following sequence can be used to trigger a UAF:

        int fscontext_fd = fsopen("cgroup");
        int fd_null = open("/dev/null, O_RDONLY);
        int fsconfig(fscontext_fd, FSCONFIG_SET_FD, "source", fd_null);
        close_range(3, ~0U, 0);

    The cgroup v1 specific fs parser expects a string for the "source"
    parameter.  However, it is perfectly legitimate to e.g.  specify a file
    descriptor for the "source" parameter.  The fs parser doesn't know what
    a filesystem allows there.  So it's a bug to assume that "source" is
    always of type fs_value_is_string when it can reasonably also be
    fs_value_is_file.

    This assumption in the cgroup code causes a UAF because struct
    fs_parameter uses a union for the actual value.  Access to that union is
    guarded by the param->type member.  Since the cgroup paramter parser
    didn't check param->type but unconditionally moved param->string into
    fc->source a close on the fscontext_fd would trigger a UAF during
    put_fs_context() which frees fc->source thereby freeing the file stashed
    in param->file causing a UAF during a close of the fd_null.

    Fix this by verifying that param->type is actually a string and report
    an error if not.

    In follow up patches I'll add a new generic helper that can be used here
    and by other filesystems instead of this error-prone copy-pasta fix.
    But fixing it in here first makes backporting a it to stable a lot
    easier.

    Fixes: 8d2451f4994f ("cgroup1: switch to option-by-option parsing")
    Reported-by: syzbot+283ce5a46486d6acdbaf@syzkaller.appspotmail.com
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Dmitry Vyukov <dvyukov@google.com>
    Cc: <stable@kernel.org>
    Cc: syzkaller-bugs <syzkaller-bugs@googlegroups.com>
    Signed-off-by: Christian Brauner <christian.brauner@ubuntu.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a6859443061ebbf6b5cc3837fc01ef3dfcac74e0
Author: Ville Syrjälä <ville.syrjala@linux.intel.com>
Date:   Wed Jun 30 19:44:13 2021 +0300

    drm/i915/gt: Fix -EDEADLK handling regression

    commit 2feeb52859fc1ab94cd35b61ada3a6ac4ff24243 upstream.

    The conversion to ww mutexes failed to address the fence code which
    already returns -EDEADLK when we run out of fences. Ww mutexes on
    the other hand treat -EDEADLK as an internal errno value indicating
```

```
    a need to restart the operation due to a deadlock. So now when the
    fence code returns -EDEADLK the higher level code erroneously
    restarts everything instead of returning the error to userspace
    as is expected.

    To remedy this let's switch the fence code to use a different errno
    value for this. -ENOBUFS seems like a semi-reasonable unique choice.
    Apart from igt the only user of this I could find is sna, and even
    there all we do is dump the current fence registers from debugfs
    into the X server log. So no user visible functionality is affected.
    If we really cared about preserving this we could of course convert
    back to -EDEADLK higher up, but doesn't seem like that's worth
    the hassle here.

    Not quite sure which commit specifically broke this, but I'll
    just attribute it to the general gem ww mutex work.

    Cc: stable@vger.kernel.org
    Cc: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    Cc: Thomas Hellström <thomas.hellstrom@intel.com>
    Testcase: igt/gem_pread/exhaustion
    Testcase: igt/gem_pwrite/basic-exhaustion
    Testcase: igt/gem_fenced_exec_thrash/too-many-fences
    Fixes: 80f0b679d6f0 ("drm/i915: Add an implementation for i915_gem_ww_ctx locking, v2.")
    Signed-off-by: Ville Syrjälä <ville.syrjala@linux.intel.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210630164413.25481-1-
ville.syrjala@linux.intel.com
    Reviewed-by: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    (cherry picked from commit 78d2ad7eb4e1f0e9cd5d79788446b6092c21d3e0)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1786384f52175794d86cd8344bfbb463c30fc231
Author: Matthew Auld <matthew.auld@intel.com>
Date:   Tue Jul 13 14:04:31 2021 +0100

    drm/i915/gtt: drop the page table optimisation

    commit 0abb33bfca0fb74df76aac03e90ce685016ef7be upstream.

    We skip filling out the pt with scratch entries if the va range covers
    the entire pt, since we later have to fill it with the PTEs for the
    object pages anyway. However this might leave open a small window where
    the PTEs don't point to anything valid for the HW to consume.

    When for example using 2M GTT pages this fill_px() showed up as being
    quite significant in perf measurements, and ends up being completely
    wasted since we ignore the pt and just use the pde directly.

    Anyway, currently we have our PTE construction split between alloc and
    insert, which is probably slightly iffy nowadays, since the alloc
    doesn't actually allocate anything anymore, instead it just sets up the
    page directories and points the PTEs at the scratch page. Later when we
    do the insert step we re-program the PTEs again. Better might be to
    squash the alloc and insert into a single step, then bringing back this
    optimisation(along with some others) should be possible.

    Fixes: 14826673247e ("drm/i915: Only initialize partially filled pagetables")
    Signed-off-by: Matthew Auld <matthew.auld@intel.com>
    Cc: Jon Bloomfield <jon.bloomfield@intel.com>
    Cc: Chris Wilson <chris.p.wilson@intel.com>
    Cc: Daniel Vetter <daniel@ffwll.ch>
    Cc: <stable@vger.kernel.org> # v4.15+
    Reviewed-by: Daniel Vetter <daniel.vetter@ffwll.ch>
    Link: https://patchwork.freedesktop.org/patch/msgid/20210713130431.2392740-1-
```

```
matthew.auld@intel.com
    (cherry picked from commit 8f88ca76b3942d82e2c1cea8735ec368d89ecc15)
    Signed-off-by: Rodrigo Vivi <rodrigo.vivi@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e67361d6c083ff3f6b43b3208cc1ea41c9729300
Author: Jinzhou Su <Jinzhou.Su@amd.com>
Date:   Tue Jul 13 09:26:11 2021 +0800

    drm/amdgpu: add another Renoir DID

    commit 775da83005cb61d4c213c636df9337da05714ff1 upstream.

    Add new PCI device id.

    Signed-off-by: Jinzhou Su <Jinzhou.Su@amd.com>
    Reviewed-by: Huang Rui <ray.huang@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Cc: stable@vger.kernel.org # 5.11.x
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 38379a09f7355a794d9a31c46542f58cb16a650e
Author: Steven Rostedt (VMware) <rostedt@goodmis.org>
Date:   Thu Jul 15 00:02:06 2021 -0400

    tracing: Do not reference char * as a string in histograms

    commit 704adfb5a9978462cd861f170201ae2b5e3d3a80 upstream.

    The histogram logic was allowing events with char * pointers to be used as
    normal strings. But it was easy to crash the kernel with:

     # echo 'hist:keys=filename' > events/syscalls/sys_enter_openat/trigger

    And open some files, and boom!

     BUG: unable to handle page fault for address: 00007f2ced0c3280
     #PF: supervisor read access in kernel mode
     #PF: error_code(0x0000) - not-present page
     PGD 1173fa067 P4D 1173fa067 PUD 1171b6067 PMD 1171dd067 PTE 0
     Oops: 0000 [#1] PREEMPT SMP
     CPU: 6 PID: 1810 Comm: cat Not tainted 5.13.0-rc5-test+ #61
     Hardware name: Hewlett-Packard HP Compaq Pro 6300 SFF/339A, BIOS K01
    v03.03 07/14/2016
     RIP: 0010:strlen+0x0/0x20
     Code: f6 82 80 2a 0b a9 20 74 11 0f b6 50 01 48 83 c0 01 f6 82 80 2a 0b
    a9 20 75 ef c3 66 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 40 00 <80> 3f 00 74
    10 48 89 f8 48 83 c0 01 80 38 00 75 f7 48 29 f8 c3

     RSP: 0018:ffffbdbf81567b50 EFLAGS: 00010246
     RAX: 0000000000000003 RBX: ffff93815cdb3800 RCX: ffff9382401a22d0
     RDX: 0000000000000100 RSI: 0000000000000000 RDI: 00007f2ced0c3280
     RBP: 0000000000000100 R08: ffff9382409ff074 R09: ffffbdbf81567c98
     R10: ffff9382409ff074 R11: 0000000000000000 R12: ffff9382409ff074
     R13: 0000000000000001 R14: ffff93815a744f00 R15: 00007f2ced0c3280
     FS:  00007f2ced0f8580(0000) GS:ffff93825a800000(0000)
    knlGS:0000000000000000
     CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
     CR2: 00007f2ced0c3280 CR3: 0000000107069005 CR4: 00000000001706e0
     Call Trace:
      event_hist_trigger+0x463/0x5f0
      ? find_held_lock+0x32/0x90
      ? sched_clock_cpu+0xe/0xd0
      ? lock_release+0x155/0x440
```

```
       ? kernel_init_free_pages+0x6d/0x90
       ? preempt_count_sub+0x9b/0xd0
       ? kernel_init_free_pages+0x6d/0x90
       ? get_page_from_freelist+0x12c4/0x1680
       ? __rb_reserve_next+0xe5/0x460
       ? ring_buffer_lock_reserve+0x12a/0x3f0
       event_triggers_call+0x52/0xe0
       ftrace_syscall_enter+0x264/0x2c0
       syscall_trace_enter.constprop.0+0x1ee/0x210
       do_syscall_64+0x1c/0x80
       entry_SYSCALL_64_after_hwframe+0x44/0xae

    Where it triggered a fault on strlen(key) where key was the filename.

    The reason is that filename is a char * to user space, and the histogram
    code just blindly dereferenced it, with obvious bad results.

    I originally tried to use strncpy_from_user/kernel_nofault() but found
    that there's other places that its dereferenced and not worth the effort.

    Just do not allow "char *" to act like strings.

    Link: https://lkml.kernel.org/r/20210715000206.025df9d2@rorschach.local.home

    Cc: Ingo Molnar <mingo@kernel.org>
    Cc: Andrew Morton <akpm@linux-foundation.org>
    Cc: Masami Hiramatsu <mhiramat@kernel.org>
    Cc: Tzvetomir Stoyanov <tz.stoyanov@gmail.com>
    Cc: stable@vger.kernel.org
    Acked-by: Namhyung Kim <namhyung@kernel.org>
    Acked-by: Tom Zanussi <zanussi@kernel.org>
    Fixes: 79e577cbce4c4 ("tracing: Support string type key properly")
    Fixes: 5967bd5c4239 ("tracing: Let filter_assign_type() detect FILTER_PTR_STRING")
    Signed-off-by: Steven Rostedt (VMware) <rostedt@goodmis.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7e6a4c304debd9be7cacecbe99b3481985e0c885
Author: Lu Baolu <baolu.lu@linux.intel.com>
Date:   Mon Jul 12 15:17:12 2021 +0800

    iommu/vt-d: Fix clearing real DMA device's scalable-mode context entries

    commit 474dd1c6506411752a9b2f2233eec11f1733a099 upstream.

    The commit 2b0140c69637e ("iommu/vt-d: Use pci_real_dma_dev() for mapping")
    fixes an issue of "sub-device is removed where the context entry is cleared
    for all aliases". But this commit didn't consider the PASID entry and PASID
    table in VT-d scalable mode. This fix increases the coverage of scalable
    mode.

    Suggested-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Fixes: 8038bdb855331 ("iommu/vt-d: Only clear real DMA device's context entries")
    Fixes: 2b0140c69637e ("iommu/vt-d: Use pci_real_dma_dev() for mapping")
    Cc: stable@vger.kernel.org # v5.6+
    Cc: Jon Derrick <jonathan.derrick@intel.com>
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Link: https://lore.kernel.org/r/20210712071712.3416949-1-baolu.lu@linux.intel.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 948ec6d003280d49aca49b366aa5cb140f87434d
Author: Sanjay Kumar <sanjay.k.kumar@intel.com>
Date:   Mon Jul 12 15:13:15 2021 +0800

    iommu/vt-d: Global devTLB flush when present context entry changed
```

```
    commit 37764b952e1b39053defc7ebe5dcd8c4e3e78de9 upstream.

    This fixes a bug in context cache clear operation. The code was not
    following the correct invalidation flow. A global device TLB invalidation
    should be added after the IOTLB invalidation. At the same time, it
    uses the domain ID from the context entry. But in scalable mode, the
    domain ID is in PASID table entry, not context entry.

    Fixes: 7373a8cc38197 ("iommu/vt-d: Setup context and enable RID2PASID support")
    Cc: stable@vger.kernel.org # v5.0+
    Signed-off-by: Sanjay Kumar <sanjay.k.kumar@intel.com>
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Link: https://lore.kernel.org/r/20210712071315.3416543-1-baolu.lu@linux.intel.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 208e33fa8359bac749851a34cd8c5dbf66f3f643
Author: Steffen Maier <maier@linux.ibm.com>
Date:   Fri Jul 2 18:09:22 2021 +0200

    scsi: zfcp: Report port fc_security as unknown early during remote cable pull

    commit 8b3bdd99c092bbaeaa7d9eecb1a3e5dc9112002b upstream.

    On remote cable pull, a zfcp_port keeps its status and only gets
    ZFCP_STATUS_PORT_LINK_TEST added. Only after an ADISC timeout, we would
    actually start port recovery and remove ZFCP_STATUS_COMMON_UNBLOCKED which
    zfcp_sysfs_port_fc_security_show() detected and reported as "unknown"
    instead of the old and possibly stale zfcp_port->connection_info.

    Add check for ZFCP_STATUS_PORT_LINK_TEST for timely "unknown" report.

    Link: https://lore.kernel.org/r/20210702160922.2667874-1-maier@linux.ibm.com
    Fixes: a17c78460093 ("scsi: zfcp: report FC Endpoint Security in sysfs")
    Cc: <stable@vger.kernel.org> #5.7+
    Reviewed-by: Benjamin Block <bblock@linux.ibm.com>
    Signed-off-by: Steffen Maier <maier@linux.ibm.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c1671d2d2ef8a84837eea1b4d99ca0c6a66fb691
Author: Tyrel Datwyler <tyreld@linux.ibm.com>
Date:   Thu Jul 1 13:56:59 2021 -0600

    scsi: core: Fix bad pointer dereference when ehandler kthread is invalid

    commit 93aa71ad7379900e61c8adff6a710a4c18c7c99b upstream.

    Commit 66a834d09293 ("scsi: core: Fix error handling of scsi_host_alloc()")
    changed the allocation logic to call put_device() to perform host cleanup
    with the assumption that IDA removal and stopping the kthread would
    properly be performed in scsi_host_dev_release(). However, in the unlikely
    case that the error handler thread fails to spawn, shost->ehandler is set
    to ERR_PTR(-ENOMEM).

    The error handler cleanup code in scsi_host_dev_release() will call
    kthread_stop() if shost->ehandler != NULL which will always be the case
    whether the kthread was successfully spawned or not. In the case that it
    failed to spawn this has the nasty side effect of trying to dereference an
    invalid pointer when kthread_stop() is called. The following splat provides
    an example of this behavior in the wild:

    scsi host11: error handler thread failed to spawn, error = -4
    Kernel attempted to read user page (10c) - exploit attempt? (uid: 0)
```

```
    BUG: Kernel NULL pointer dereference on read at 0x0000010c
    Faulting instruction address: 0xc00000000818e9a8
    Oops: Kernel access of bad area, sig: 11 [#1]
    LE PAGE_SIZE=64K MMU=Hash SMP NR_CPUS=2048 NUMA pSeries
    Modules linked in: ibmvscsi(+) scsi_transport_srp dm_multipath dm_mirror dm_region
     hash dm_log dm_mod fuse overlay squashfs loop
    CPU: 12 PID: 274 Comm: systemd-udevd Not tainted 5.13.0-rc7 #1
    NIP:  c00000000818e9a8 LR: c0000000089846e8 CTR: 0000000000007ee8
    REGS: c000000037d12ea0 TRAP: 0300   Not tainted  (5.13.0-rc7)
    MSR:  800000000280b033 &lt;SF,VEC,VSX,EE,FP,ME,IR,DR,RI,LE&gt;  CR: 28228228
    XER: 20040001
    CFAR: c0000000089846e4 DAR: 000000000000010c DSISR: 40000000 IRQMASK: 0
    GPR00: c0000000089846e8 c000000037d13140 c000000009cc1100 fffffffffffffffc
    GPR04: 0000000000000001 0000000000000000 0000000000000000 c000000037dc0000
    GPR08: 0000000000000000 c000000037dc0000 0000000000000001 00000000fffff7ff
    GPR12: 0000000000008000 c00000000a049000 c000000037d13d00 000000011134d5a0
    GPR16: 0000000000001740 c0080000190d0000 c0080000190d1740 c000000009129288
    GPR20: c000000037d13bc0 0000000000000001 c000000037d13bc0 c0080000190b7898
    GPR24: c0080000190b7708 0000000000000000 c000000033bb2c48 0000000000000000
    GPR28: c000000046b28280 0000000000000000 000000000000010c fffffffffffffffc
    NIP [c00000000818e9a8] kthread_stop+0x38/0x230
    LR [c0000000089846e8] scsi_host_dev_release+0x98/0x160
    Call Trace:
    [c000000033bb2c48] 0xc000000033bb2c48 (unreliable)
    [c0000000089846e8] scsi_host_dev_release+0x98/0x160
    [c00000000891e960] device_release+0x60/0x100
    [c0000000087e55c4] kobject_release+0x84/0x210
    [c00000000891ec78] put_device+0x28/0x40
    [c000000008984ea4] scsi_host_alloc+0x314/0x430
    [c0080000190b38bc] ibmvscsi_probe+0x54/0xad0 [ibmvscsi]
    [c000000008110104] vio_bus_probe+0xa4/0x4b0
    [c00000000892a860] really_probe+0x140/0x680
    [c00000000892aefc] driver_probe_device+0x15c/0x200
    [c00000000892b63c] device_driver_attach+0xcc/0xe0
    [c00000000892b740] __driver_attach+0xf0/0x200
    [c000000008926f28] bus_for_each_dev+0xa8/0x130
    [c000000008929ce4] driver_attach+0x34/0x50
    [c000000008928fc0] bus_add_driver+0x1b0/0x300
    [c00000000892c798] driver_register+0x98/0x1a0
    [c00000000810eb60] __vio_register_driver+0x80/0xe0
    [c0080000190b4a30] ibmvscsi_module_init+0x9c/0xdc [ibmvscsi]
    [c0000000080121d0] do_one_initcall+0x60/0x2d0
    [c000000008261abc] do_init_module+0x7c/0x320
    [c000000008265700] load_module+0x2350/0x25b0
    [c000000008265cb4] __do_sys_finit_module+0xd4/0x160
    [c000000008031110] system_call_exception+0x150/0x2d0
    [c00000000800d35c] system_call_common+0xec/0x278

    Fix this be nulling shost->ehandler when the kthread fails to spawn.

    Link: https://lore.kernel.org/r/20210701195659.3185475-1-tyreld@linux.ibm.com
    Fixes: 66a834d09293 ("scsi: core: Fix error handling of scsi_host_alloc()")
    Cc: stable@vger.kernel.org
    Reviewed-by: Ming Lei <ming.lei@redhat.com>
    Signed-off-by: Tyrel Datwyler <tyreld@linux.ibm.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit aa25872fcc35b94f7684d0305440ce5e53bf69d3
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Wed Jul 7 15:50:59 2021 +0300

    KVM: SVM: remove INIT intercept handler

    commit 896707c212d440a6863ce0a3930c8a609e24497d upstream.
```

```
    Kernel never sends real INIT even to CPUs, other than on boot.

    Thus INIT interception is an error which should be caught
    by a check for an unknown VMexit reason.

    On top of that, the current INIT VM exit handler skips
    the current instruction which is wrong.
    That was added in commit 5ff3a351f687 ("KVM: x86: Move trivial
    instruction-based exit handlers to common code").

    Fixes: 5ff3a351f687 ("KVM: x86: Move trivial instruction-based exit handlers to common
code")
    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Message-Id: <20210707125100.677203-3-mlevitsk@redhat.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bbca8ad77f260b6d7f468b26e819ae7f282b2622
Author: Maxim Levitsky <mlevitsk@redhat.com>
Date:   Wed Jul 7 15:50:58 2021 +0300

    KVM: SVM: #SMI interception must not skip the instruction

    commit 991afbbee8ac93b055a27477278a5fb556af1ff4 upstream.

    Commit 5ff3a351f687 ("KVM: x86: Move trivial instruction-based
    exit handlers to common code"), unfortunately made a mistake of
    treating nop_on_interception and nop_interception in the same way.

    Former does truly nothing while the latter skips the instruction.

    SMI VM exit handler should do nothing.
    (SMI itself is handled by the host when we do STGI)

    Fixes: 5ff3a351f687 ("KVM: x86: Move trivial instruction-based exit handlers to common
code")
    Signed-off-by: Maxim Levitsky <mlevitsk@redhat.com>
    Message-Id: <20210707125100.677203-2-mlevitsk@redhat.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2338682c3fd14f4a718da80ceffdb33ecedbf072
Author: Lai Jiangshan <laijs@linux.alibaba.com>
Date:   Tue Jun 29 01:26:32 2021 +0800

    KVM: X86: Disable hardware breakpoints unconditionally before kvm_x86->run()

    commit f85d40160691881a17a397c448d799dfc90987ba upstream.

    When the host is using debug registers but the guest is not using them
    nor is the guest in guest-debug state, the kvm code does not reset
    the host debug registers before kvm_x86->run().  Rather, it relies on
    the hardware vmentry instruction to automatically reset the dr7 registers
    which ensures that the host breakpoints do not affect the guest.

    This however violates the non-instrumentable nature around VM entry
    and exit; for example, when a host breakpoint is set on vcpu->arch.cr2,

    Another issue is consistency.  When the guest debug registers are active,
    the host breakpoints are reset before kvm_x86->run(). But when the
    guest debug registers are inactive, the host breakpoints are delayed to
    be disabled.  The host tracing tools may see different results depending
```

```
    on what the guest is doing.

    To fix the problems, we clear %db7 unconditionally before kvm_x86->run()
    if the host has set any breakpoints, no matter if the guest is using
    them or not.

    Signed-off-by: Lai Jiangshan <laijs@linux.alibaba.com>
    Message-Id: <20210628172632.81029-1-jiangshanlai@gmail.com>
    Cc: stable@vger.kernel.org
    [Only clear %db7 instead of reloading all debug registers. - Paolo]
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f4e862a6a3d950e1e6196e1d7eebfc21d66533c0
Author: Vitaly Kuznetsov <vkuznets@redhat.com>
Date:   Mon Jun 28 12:44:20 2021 +0200

    KVM: nSVM: Check the value written to MSR_VM_HSAVE_PA

    commit fce7e152ffc8f89d02a80617b16c7aa1527847c8 upstream.

    APM states that #GP is raised upon write to MSR_VM_HSAVE_PA when
    the supplied address is not page-aligned or is outside of "maximum
    supported physical address for this implementation".
    page_address_valid() check seems suitable. Also, forcefully page-align
    the address when it's written from VMM.

    Signed-off-by: Vitaly Kuznetsov <vkuznets@redhat.com>
    Message-Id: <20210628104425.391276-2-vkuznets@redhat.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Maxim Levitsky <mlevitsk@redhat.com>
    [Add comment about behavior for host-provided values. - Paolo]
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a238f763e6502d3042b9a5ebbd41cbc23e5e1527
Author: Sean Christopherson <seanjc@google.com>
Date:   Thu Jun 24 19:03:54 2021 -0700

    KVM: SVM: Revert clearing of C-bit on GPA in #NPF handler

    commit 76ff371b67cb12fb635396234468abcf6a466f16 upstream.

    Don't clear the C-bit in the #NPF handler, as it is a legal GPA bit for
    non-SEV guests, and for SEV guests the C-bit is dropped before the GPA
    hits the NPT in hardware.  Clearing the bit for non-SEV guests causes KVM
    to mishandle #NPFs with that collide with the host's C-bit.

    Although the APM doesn't explicitly state that the C-bit is not reserved
    for non-SEV, Tom Lendacky confirmed that the following snippet about the
    effective reduction due to the C-bit does indeed apply only to SEV guests.

      Note that because guest physical addresses are always translated
      through the nested page tables, the size of the guest physical address
      space is not impacted by any physical address space reduction indicated
      in CPUID 8000_001F[EBX]. If the C-bit is a physical address bit however,
      the guest physical address space is effectively reduced by 1 bit.

    And for SEV guests, the APM clearly states that the bit is dropped before
    walking the nested page tables.

      If the C-bit is an address bit, this bit is masked from the guest
      physical address when it is translated through the nested page tables.
      Consequently, the hypervisor does not need to be aware of which pages
      the guest has chosen to mark private.
```

```
    Note, the bogus C-bit clearing was removed from legacy #PF handler in
    commit 6d1b867d0456 ("KVM: SVM: Don't strip the C-bit from CR2 on #PF
    interception").

    Fixes: 0ede79e13224 ("KVM: SVM: Clear C-bit from the page fault address")
    Cc: Peter Gonda <pgonda@google.com>
    Cc: Brijesh Singh <brijesh.singh@amd.com>
    Cc: Tom Lendacky <thomas.lendacky@amd.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20210625020354.431829-3-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 555e674fd3f9088f6a610e1d3ed72c2230ce6094
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Jun 23 16:05:49 2021 -0700

    KVM: x86/mmu: Do not apply HPA (memory encryption) mask to GPAs

    commit fc9bf2e087efcd81bda2e52d09616d2a1bf982a8 upstream.

    Ignore "dynamic" host adjustments to the physical address mask when
    generating the masks for guest PTEs, i.e. the guest PA masks.  The host
    physical address space and guest physical address space are two different
    beasts, e.g. even though SEV's C-bit is the same bit location for both
    host and guest, disabling SME in the host (which clears shadow_me_mask)
    does not affect the guest PTE->GPA "translation".

    For non-SEV guests, not dropping bits is the correct behavior.  Assuming
    KVM and userspace correctly enumerate/configure guest MAXPHYADDR, bits
    that are lost as collateral damage from memory encryption are treated as
    reserved bits, i.e. KVM will never get to the point where it attempts to
    generate a gfn using the affected bits.  And if userspace wants to create
    a bogus vCPU, then userspace gets to deal with the fallout of hardware
    doing odd things with bad GPAs.

    For SEV guests, not dropping the C-bit is technically wrong, but it's a
    moot point because KVM can't read SEV guest's page tables in any case
    since they're always encrypted.  Not to mention that the current KVM code
    is also broken since sme_me_mask does not have to be non-zero for SEV to
    be supported by KVM.  The proper fix would be to teach all of KVM to
    correctly handle guest private memory, but that's a task for the future.

    Fixes: d0ec49d4de90 ("kvm/x86/svm: Support Secure Memory Encryption within KVM")
    Cc: stable@vger.kernel.org
    Cc: Brijesh Singh <brijesh.singh@amd.com>
    Cc: Tom Lendacky <thomas.lendacky@amd.com>
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20210623230552.4027702-5-seanjc@google.com>
    [Use a new header instead of adding header guards to paging_tmpl.h. - Paolo]
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4ccf1f873b7029467697f7ae9df6bf6625da1cab
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Jun 23 16:05:47 2021 -0700

    KVM: x86: Use kernel's x86_phys_bits to handle reduced MAXPHYADDR

    commit e39f00f60ebd2e7b295c37a05e6349df656d3eb8 upstream.

    Use boot_cpu_data.x86_phys_bits instead of the raw CPUID information to
    enumerate the MAXPHYADDR for KVM guests when TDP is disabled (the guest
```

```
    version is only relevant to NPT/TDP).

    When using shadow paging, any reductions to the host's MAXPHYADDR apply
    to KVM and its guests as well, i.e. using the raw CPUID info will cause
    KVM to misreport the number of PA bits available to the guest.

    Unconditionally zero out the "Physical Address bit reduction" entry.
    For !TDP, the adjustment is already done, and for TDP enumerating the
    host's reduction is wrong as the reduction does not apply to GPAs.

    Fixes: 9af9b94068fb ("x86/cpu/AMD: Handle SME reduction in physical address size")
    Cc: stable@vger.kernel.org
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20210623230552.4027702-3-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0eb28e18e56cb407d15dc4becac03720d2f6c104
Author: Sean Christopherson <seanjc@google.com>
Date:   Wed Jun 23 16:05:46 2021 -0700

    KVM: x86: Use guest MAXPHYADDR from CPUID.0x8000_0008 iff TDP is enabled

    commit 4bf48e3c0aafd32b960d341c4925b48f416f14a5 upstream.

    Ignore the guest MAXPHYADDR reported by CPUID.0x8000_0008 if TDP, i.e.
    NPT, is disabled, and instead use the host's MAXPHYADDR.  Per AMD'S APM:

      Maximum guest physical address size in bits. This number applies only
      to guests using nested paging. When this field is zero, refer to the
      PhysAddrSize field for the maximum guest physical address size.

    Fixes: 24c82e576b78 ("KVM: Sanitize cpuid")
    Cc: stable@vger.kernel.org
    Signed-off-by: Sean Christopherson <seanjc@google.com>
    Message-Id: <20210623230552.4027702-2-seanjc@google.com>
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 228ce0ed1283000073e25fe6196aed1fb879467e
Author: Christian Borntraeger <borntraeger@de.ibm.com>
Date:   Thu Jul 1 17:42:24 2021 +0200

    KVM: selftests: do not require 64GB in set_memory_region_test

    commit cd4220d23bf3f43cf720e82bdee681f383433ae2 upstream.

    Unless the user sets overcommit_memory or has plenty of swap, the latest
    changes to the testcase will result in ENOMEM failures for hosts with
    less than 64GB RAM. As we do not use much of the allocated memory, we
    can use MAP_NORESERVE to avoid this error.

    Cc: Zenghui Yu <yuzenghui@huawei.com>
    Cc: vkuznets@redhat.com
    Cc: wanghaibin.wang@huawei.com
    Cc: stable@vger.kernel.org
    Fixes: 309505dd5685 ("KVM: selftests: Fix mapping length truncation in m{,un}map()")
    Tested-by: Zenghui Yu <yuzenghui@huawei.com>
    Link: https://lore.kernel.org/kvm/20210701160425.33666-1-borntraeger@de.ibm.com/
    Signed-off-by: Christian Borntraeger <borntraeger@de.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 069d44a24c0ff8f85adf49233aae7a8ca16f5c7e
Author: Kefeng Wang <wangkefeng.wang@huawei.com>
Date:   Sat Jun 26 15:03:04 2021 +0800
```

```
    KVM: mmio: Fix use-after-free Read in kvm_vm_ioctl_unregister_coalesced_mmio

    commit 23fa2e46a5556f787ce2ea1a315d3ab93cced204 upstream.

    BUG: KASAN: use-after-free in kvm_vm_ioctl_unregister_coalesced_mmio+0x7c/0x1ec
arch/arm64/kvm/../../../virt/kvm/coalesced_mmio.c:183
    Read of size 8 at addr ffff0000c03a2500 by task syz-executor083/4269

    CPU: 5 PID: 4269 Comm: syz-executor083 Not tainted 5.10.0 #7
    Hardware name: linux,dummy-virt (DT)
    Call trace:
     dump_backtrace+0x0/0x2d0 arch/arm64/kernel/stacktrace.c:132
     show_stack+0x28/0x34 arch/arm64/kernel/stacktrace.c:196
     __dump_stack lib/dump_stack.c:77 [inline]
     dump_stack+0x110/0x164 lib/dump_stack.c:118
     print_address_description+0x78/0x5c8 mm/kasan/report.c:385
     __kasan_report mm/kasan/report.c:545 [inline]
     kasan_report+0x148/0x1e4 mm/kasan/report.c:562
     check_memory_region_inline mm/kasan/generic.c:183 [inline]
     __asan_load8+0xb4/0xbc mm/kasan/generic.c:252
     kvm_vm_ioctl_unregister_coalesced_mmio+0x7c/0x1ec
arch/arm64/kvm/../../../virt/kvm/coalesced_mmio.c:183
     kvm_vm_ioctl+0xe30/0x14c4 arch/arm64/kvm/../../../virt/kvm/kvm_main.c:3755
     vfs_ioctl fs/ioctl.c:48 [inline]
     __do_sys_ioctl fs/ioctl.c:753 [inline]
     __se_sys_ioctl fs/ioctl.c:739 [inline]
     __arm64_sys_ioctl+0xf88/0x131c fs/ioctl.c:739
     __invoke_syscall arch/arm64/kernel/syscall.c:36 [inline]
     invoke_syscall arch/arm64/kernel/syscall.c:48 [inline]
     el0_svc_common arch/arm64/kernel/syscall.c:158 [inline]
     do_el0_svc+0x120/0x290 arch/arm64/kernel/syscall.c:220
     el0_svc+0x1c/0x28 arch/arm64/kernel/entry-common.c:367
     el0_sync_handler+0x98/0x170 arch/arm64/kernel/entry-common.c:383
     el0_sync+0x140/0x180 arch/arm64/kernel/entry.S:670

    Allocated by task 4269:
     stack_trace_save+0x80/0xb8 kernel/stacktrace.c:121
     kasan_save_stack mm/kasan/common.c:48 [inline]
     kasan_set_track mm/kasan/common.c:56 [inline]
     __kasan_kmalloc+0xdc/0x120 mm/kasan/common.c:461
     kasan_kmalloc+0xc/0x14 mm/kasan/common.c:475
     kmem_cache_alloc_trace include/linux/slab.h:450 [inline]
     kmalloc include/linux/slab.h:552 [inline]
     kzalloc include/linux/slab.h:664 [inline]
     kvm_vm_ioctl_register_coalesced_mmio+0x78/0x1cc
arch/arm64/kvm/../../../virt/kvm/coalesced_mmio.c:146
     kvm_vm_ioctl+0x7e8/0x14c4 arch/arm64/kvm/../../../virt/kvm/kvm_main.c:3746
     vfs_ioctl fs/ioctl.c:48 [inline]
     __do_sys_ioctl fs/ioctl.c:753 [inline]
     __se_sys_ioctl fs/ioctl.c:739 [inline]
     __arm64_sys_ioctl+0xf88/0x131c fs/ioctl.c:739
     __invoke_syscall arch/arm64/kernel/syscall.c:36 [inline]
     invoke_syscall arch/arm64/kernel/syscall.c:48 [inline]
     el0_svc_common arch/arm64/kernel/syscall.c:158 [inline]
     do_el0_svc+0x120/0x290 arch/arm64/kernel/syscall.c:220
     el0_svc+0x1c/0x28 arch/arm64/kernel/entry-common.c:367
     el0_sync_handler+0x98/0x170 arch/arm64/kernel/entry-common.c:383
     el0_sync+0x140/0x180 arch/arm64/kernel/entry.S:670

    Freed by task 4269:
     stack_trace_save+0x80/0xb8 kernel/stacktrace.c:121
     kasan_save_stack mm/kasan/common.c:48 [inline]
     kasan_set_track+0x38/0x6c mm/kasan/common.c:56
     kasan_set_free_info+0x20/0x40 mm/kasan/generic.c:355
```

```
      __kasan_slab_free+0x124/0x150 mm/kasan/common.c:422
      kasan_slab_free+0x10/0x1c mm/kasan/common.c:431
      slab_free_hook mm/slub.c:1544 [inline]
      slab_free_freelist_hook mm/slub.c:1577 [inline]
      slab_free mm/slub.c:3142 [inline]
      kfree+0x104/0x38c mm/slub.c:4124
      coalesced_mmio_destructor+0x94/0xa4
 arch/arm64/kvm/../../../virt/kvm/coalesced_mmio.c:102
      kvm_iodevice_destructor include/kvm/iodev.h:61 [inline]
      kvm_io_bus_unregister_dev+0x248/0x280 arch/arm64/kvm/../../../virt/kvm/kvm_main.c:4374
      kvm_vm_ioctl_unregister_coalesced_mmio+0x158/0x1ec
 arch/arm64/kvm/../../../virt/kvm/coalesced_mmio.c:186
      kvm_vm_ioctl+0xe30/0x14c4 arch/arm64/kvm/../../../virt/kvm/kvm_main.c:3755
      vfs_ioctl fs/ioctl.c:48 [inline]
      __do_sys_ioctl fs/ioctl.c:753 [inline]
      __se_sys_ioctl fs/ioctl.c:739 [inline]
      __arm64_sys_ioctl+0xf88/0x131c fs/ioctl.c:739
      __invoke_syscall arch/arm64/kernel/syscall.c:36 [inline]
      invoke_syscall arch/arm64/kernel/syscall.c:48 [inline]
      el0_svc_common arch/arm64/kernel/syscall.c:158 [inline]
      do_el0_svc+0x120/0x290 arch/arm64/kernel/syscall.c:220
      el0_svc+0x1c/0x28 arch/arm64/kernel/entry-common.c:367
      el0_sync_handler+0x98/0x170 arch/arm64/kernel/entry-common.c:383
      el0_sync+0x140/0x180 arch/arm64/kernel/entry.S:670

    If kvm_io_bus_unregister_dev() return -ENOMEM, we already call kvm_iodevice_destructor()
    inside this function to delete 'struct kvm_coalesced_mmio_dev *dev' from list
    and free the dev, but kvm_iodevice_destructor() is called again, it will lead
    the above issue.

    Let's check the the return value of kvm_io_bus_unregister_dev(), only call
    kvm_iodevice_destructor() if the return value is 0.

    Cc: Paolo Bonzini <pbonzini@redhat.com>
    Cc: kvm@vger.kernel.org
    Reported-by: Hulk Robot <hulkci@huawei.com>
    Signed-off-by: Kefeng Wang <wangkefeng.wang@huawei.com>
    Message-Id: <20210626070304.143456-1-wangkefeng.wang@huawei.com>
    Cc: stable@vger.kernel.org
    Fixes: 5d3c4c79384a ("KVM: Stop looking for coalesced MMIO zones if the bus is
 destroyed", 2021-04-20)
    Signed-off-by: Paolo Bonzini <pbonzini@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

 commit d4e376c83736c49d2150b2d8cb13c5436083370f
 Author: Ronnie Sahlberg <lsahlber@redhat.com>
 Date:   Tue Jul 13 12:22:59 2021 +1000

    cifs: Do not use the original cruid when following DFS links for multiuser mounts

    commit 50630b3f1ada0bf412d3f28e73bac310448d9d6f upstream.

    Bugzilla: https://bugzilla.kernel.org/show_bug.cgi?id=213565

    cruid should only be used for the initial mount and after this we should use the current
    users credentials.
    Ignore the original cruid mount argument when creating a new context for a multiuser
 mount
    following a DFS link.

    Fixes: 24e0a1eff9e2 ("cifs: switch to new mount api")
    Cc: stable@vger.kernel.org # 5.11+
    Reported-by: Xiaoli Feng <xifeng@redhat.com>
    Signed-off-by: Ronnie Sahlberg <lsahlber@redhat.com>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
```

```
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0c480c45bed33354316e8c4f19a3cf53c701732a
Author: Paulo Alcantara <pc@cjr.nz>
Date:   Mon Jul 12 12:38:24 2021 -0300

    cifs: handle reconnect of tcon when there is no cached dfs referral

    commit 507345b5ae6a57b7ecd7550ff39282ed20de7b8d upstream.

    When there is no cached DFS referral of tcon->dfs_path, then reconnect
    to same share.

    Signed-off-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a29b76ec20f2337557801243cb9fa76bc62155b7
Author: Shyam Prasad N <sprasad@microsoft.com>
Date:   Tue May 18 15:05:50 2021 +0000

    cifs: use the expiry output of dns_query to schedule next resolution

    commit 506c1da44fee32ba1d3a70413289ad58c772bba6 upstream.

    We recently fixed DNS resolution of the server hostname during reconnect.
    However, server IP address may change, even when the old one continues
    to server (although sub-optimally).

    We should schedule the next DNS resolution based on the TTL of
    the DNS record used for the last resolution. This way, we resolve the
    server hostname again when a DNS record expires.

    Signed-off-by: Shyam Prasad N <sprasad@microsoft.com>
    Reviewed-by: Paulo Alcantara (SUSE) <pc@cjr.nz>
    Cc: <stable@vger.kernel.org> # v5.11+
    Signed-off-by: Steve French <stfrench@microsoft.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```