```
commit 238589d0f7b421aae18c5704dc931595019fa6c7
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Wed Sep 13 09:53:58 2023 +0200

    Linux 6.5.3

    Link: https://lore.kernel.org/r/20230911134650.921299741@linuxfoundation.org
    Tested-by: Ronald Warsow <rwarsow@gmx.de>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Bagas Sanjaya <bagasdotme@gmail.com>
    Tested-by: Linux Kernel Functional Testing <lkft@linaro.org>
    Tested-by: Sudip Mukherjee <sudip.mukherjee@codethink.co.uk>
    Tested-by: Justin M. Forbes <jforbes@fedoraproject.org>
    Tested-by: Ron Economos <re@w6rz.net>
    Tested-by: Jon Hunter <jonathanh@nvidia.com>
    Tested-by: Allen Pais <apais@linux.microsoft.com>
    Tested-by: Guenter Roeck <linux@roeck-us.net>
    Tested-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Tested-by: Salvatore Bonaccorso <carnil@debian.org>
    Tested-by: Conor Dooley <conor.dooley@microchip.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7aa2f0f86b154ce90267d3f1799b28ee410a57d3
Author: Wesley Chalmers <Wesley.Chalmers@amd.com>
Date:   Wed May 31 13:29:34 2023 -0400

    drm/amd/display: Block optimize on consecutive FAMS enables

    commit 3b6df06f01cdbff3b610b492ad4879691afdc70d upstream.

    [WHY]
    It is possible to commit state multiple times in rapid succession with
    FAMS enabled; if each of these commits were to set optimized_required,
    then the user may see latency.

    [HOW]
    fw_based_mclk_switching is currently not used in dc->clk_mgr; use it
    to track whether the current state has FAMS enabled;
    if it has, then do not disable FAMS in prepare_bandwidth, and do not set
    optimized_required.

    Reviewed-by: Rodrigo Siqueira <Rodrigo.Siqueira@amd.com>
    Signed-off-by: Wesley Chalmers <Wesley.Chalmers@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Acked-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 72260f7610480e7f9a6aa0a3ecbd59dd22147af2
Author: Andrew Morton <akpm@linux-foundation.org>
Date:   Sat Sep 2 15:59:31 2023 -0700

    revert "memfd: improve userspace warnings for missing exec-related flags".

    commit 2562d67b1bdf91c7395b0225d60fdeb26b4bc5a0 upstream.

    This warning is telling userspace developers to pass MFD_EXEC and
    MFD_NOEXEC_SEAL to memfd_create().  Commit 434ed3350f57 ("memfd: improve
    userspace warnings for missing exec-related flags") made the warning more
    frequent and visible in the hope that this would accelerate the fixing of
    errant userspace.

    But the overall effect is to generate far too much dmesg noise.
```

```
    Fixes: 434ed3350f57 ("memfd: improve userspace warnings for missing exec-related flags")
    Reported-by: Damian Tometzki <dtometzki@fedoraproject.org>
    Closes: https://lkml.kernel.org/r/ZPFzCSIgZ4QuHsSC@fedora.fritz.box
    Cc: Aleksa Sarai <cyphar@cyphar.com>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Jeff Xu <jeffxu@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a70994b2b23c1aafa12c7ed192396be2d5bc6f48
Author: Aleksa Sarai <cyphar@cyphar.com>
Date:   Mon Aug 14 18:40:59 2023 +1000

    memfd: improve userspace warnings for missing exec-related flags

    [ Upstream commit 434ed3350f57c03a9654fe0619755cc137a58935 ]

    In order to incentivise userspace to switch to passing MFD_EXEC and
    MFD_NOEXEC_SEAL, we need to provide a warning on each attempt to call
    memfd_create() without the new flags.  pr_warn_once() is not useful
    because on most systems the one warning is burned up during the boot
    process (on my system, systemd does this within the first second of boot)
    and thus userspace will in practice never see the warnings to push them to
    switch to the new flags.

    The original patchset[1] used pr_warn_ratelimited(), however there were
    concerns about the degree of spam in the kernel log[2,3].  The resulting
    inability to detect every case was flagged as an issue at the time[4].

    While we could come up with an alternative rate-limiting scheme such as
    only outputting the message if vm.memfd_noexec has been modified, or only
    outputting the message once for a given task, these alternatives have
    downsides that don't make sense given how low-stakes a single kernel
    warning message is.  Switching to pr_info_ratelimited() instead should be
    fine -- it's possible some monitoring tool will be unhappy with a stream
    of warning-level messages but there's already plenty of info-level message
    spam in dmesg.

    [1]: https://lore.kernel.org/20221215001205.51969-4-jeffxu@google.com/
    [2]: https://lore.kernel.org/202212161233.85C9783FB@keescook/
    [3]: https://lore.kernel.org/Y5yS8wCnuYGLHMj4@x1n/
    [4]: https://lore.kernel.org/f185bb42-b29c-977e-312e-3349eea15383@linuxfoundation.org/

    Link: https://lkml.kernel.org/r/20230814-memfd-vm-noexec-uapi-fixes-v2-3-
7ff9e3e10ba6@cyphar.com
    Fixes: 105ff5339f49 ("mm/memfd: add MFD_NOEXEC_SEAL and MFD_EXEC")
    Signed-off-by: Aleksa Sarai <cyphar@cyphar.com>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Dominique Martinet <asmadeus@codewreck.org>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a39eb58a57848570d0ed0ab715c64278b42c8aed
Author: Aleksa Sarai <cyphar@cyphar.com>
Date:   Mon Aug 14 18:41:00 2023 +1000
```

```
    memfd: replace ratcheting feature from vm.memfd_noexec with hierarchy

    [ Upstream commit 9876cfe8ec1cb3c88de31f4d58d57b0e7e22bcc4 ]

    This sysctl has the very unusual behaviour of not allowing any user (even
    CAP_SYS_ADMIN) to reduce the restriction setting, meaning that if you were
    to set this sysctl to a more restrictive option in the host pidns you
    would need to reboot your machine in order to reset it.

    The justification given in [1] is that this is a security feature and thus
    it should not be possible to disable.  Aside from the fact that we have
    plenty of security-related sysctls that can be disabled after being
    enabled (fs.protected_symlinks for instance), the protection provided by
    the sysctl is to stop users from being able to create a binary and then
    execute it.  A user with CAP_SYS_ADMIN can trivially do this without
    memfd_create(2):

      % cat mount-memfd.c
      #include <fcntl.h>
      #include <string.h>
      #include <stdio.h>
      #include <stdlib.h>
      #include <unistd.h>
      #include <linux/mount.h>

      #define SHELLCODE "#!/bin/echo this file was executed from this totally private
 tmpfs:"

      int main(void)
      {
            int fsfd = fsopen("tmpfs", FSOPEN_CLOEXEC);
            assert(fsfd >= 0);
            assert(!fsconfig(fsfd, FSCONFIG_CMD_CREATE, NULL, NULL, 2));

            int dfd = fsmount(fsfd, FSMOUNT_CLOEXEC, 0);
            assert(dfd >= 0);

            int execfd = openat(dfd, "exe", O_CREAT | O_RDWR | O_CLOEXEC, 0782);
            assert(execfd >= 0);
            assert(write(execfd, SHELLCODE, strlen(SHELLCODE)) == strlen(SHELLCODE));
            assert(!close(execfd));

            char *execpath = NULL;
            char *argv[] = { "bad-exe", NULL }, *envp[] = { NULL };
            execfd = openat(dfd, "exe", O_PATH | O_CLOEXEC);
            assert(execfd >= 0);
            assert(asprintf(&execpath, "/proc/self/fd/%d", execfd) > 0);
            assert(!execve(execpath, argv, envp));
      }
      % ./mount-memfd
      this file was executed from this totally private tmpfs: /proc/self/fd/5
      %

    Given that it is possible for CAP_SYS_ADMIN users to create executable
    binaries without memfd_create(2) and without touching the host filesystem
    (not to mention the many other things a CAP_SYS_ADMIN process would be
    able to do that would be equivalent or worse), it seems strange to cause a
    fair amount of headache to admins when there doesn't appear to be an
    actual security benefit to blocking this.  There appear to be concerns
    about confused-deputy-esque attacks[2] but a confused deputy that can
    write to arbitrary sysctls is a bigger security issue than executable
    memfds.

    /* New API */
```

```
    The primary requirement from the original author appears to be more based
    on the need to be able to restrict an entire system in a hierarchical
    manner[3], such that child namespaces cannot re-enable executable memfds.

    So, implement that behaviour explicitly -- the vm.memfd_noexec scope is
    evaluated up the pidns tree to &init_pid_ns and you have the most
    restrictive value applied to you.  The new lower limit you can set
    vm.memfd_noexec is whatever limit applies to your parent.

    Note that a pidns will inherit a copy of the parent pidns's effective
    vm.memfd_noexec setting at unshare() time.  This matches the existing
    behaviour, and it also ensures that a pidns will never have its
    vm.memfd_noexec setting *lowered* behind its back (but it will be raised
    if the parent raises theirs).

    /* Backwards Compatibility */

    As the previous version of the sysctl didn't allow you to lower the
    setting at all, there are no backwards compatibility issues with this
    aspect of the change.

    However it should be noted that now that the setting is completely
    hierarchical.  Previously, a cloned pidns would just copy the current
    pidns setting, meaning that if the parent's vm.memfd_noexec was changed it
    wouldn't propoagate to existing pid namespaces.  Now, the restriction
    applies recursively.  This is a uAPI change, however:

     * The sysctl is very new, having been merged in 6.3.
     * Several aspects of the sysctl were broken up until this patchset and
       the other patchset by Jeff Xu last month.

    And thus it seems incredibly unlikely that any real users would run into
    this issue. In the worst case, if this causes userspace isues we could
    make it so that modifying the setting follows the hierarchical rules but
    the restriction checking uses the cached copy.

    [1]:
https://lore.kernel.org/CABi2SkWnAgHK1i6iqSqPMYuNEhtHBkO8jUuCvmG3RmUB5TKHJw@mail.gmail.com/
    [2]:
https://lore.kernel.org/CALmYWFs_dNCzw_pW1yRAo4bGCPEtykroEQaowNULp7svwMLjOg@mail.gmail.com/
    [3]: https://lore.kernel.org/CALmYWFuahdUF7cT4cm7_TGLqPanuHXJ-
hVSfZt7vpTnc18DPrw@mail.gmail.com/

    Link: https://lkml.kernel.org/r/20230814-memfd-vm-noexec-uapi-fixes-v2-4-
7ff9e3e10ba6@cyphar.com
    Fixes: 105ff5339f49 ("mm/memfd: add MFD_NOEXEC_SEAL and MFD_EXEC")
    Signed-off-by: Aleksa Sarai <cyphar@cyphar.com>
    Cc: Dominique Martinet <asmadeus@codewreck.org>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Jeff Xu <jeffxu@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d5a597636caa749dcaaae841a17c12325b9f56e4
Author: Aleksa Sarai <cyphar@cyphar.com>
Date:   Mon Aug 14 18:40:58 2023 +1000

    memfd: do not -EACCES old memfd_create() users with vm.memfd_noexec=2

    [ Upstream commit 202e14222fadb246dfdf182e67de1518e86a1e20 ]
```

```
    Given the difficulty of auditing all of userspace to figure out whether
    every memfd_create() user has switched to passing MFD_EXEC and
    MFD_NOEXEC_SEAL flags, it seems far less distruptive to make it possible
    for older programs that don't make use of executable memfds to run under
    vm.memfd_noexec=2.  Otherwise, a small dependency change can result in
    spurious errors.  For programs that don't use executable memfds, passing
    MFD_NOEXEC_SEAL is functionally a no-op and thus having the same

    In addition, every failure under vm.memfd_noexec=2 needs to print to the
    kernel log so that userspace can figure out where the error came from.
    The concerns about pr_warn_ratelimited() spam that caused the switch to
    pr_warn_once()[1,2] do not apply to the vm.memfd_noexec=2 case.

    This is a user-visible API change, but as it allows programs to do
    something that would be blocked before, and the sysctl itself was broken
    and recently released, it seems unlikely this will cause any issues.

    [1]: https://lore.kernel.org/Y5yS8wCnuYGLHMj4@x1n/
    [2]: https://lore.kernel.org/202212161233.85C9783FB@keescook/

    Link: https://lkml.kernel.org/r/20230814-memfd-vm-noexec-uapi-fixes-v2-2-
7ff9e3e10ba6@cyphar.com
    Fixes: 105ff5339f49 ("mm/memfd: add MFD_NOEXEC_SEAL and MFD_EXEC")
    Signed-off-by: Aleksa Sarai <cyphar@cyphar.com>
    Cc: Dominique Martinet <asmadeus@codewreck.org>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Jeff Xu <jeffxu@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <shuah@kernel.org>
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 972dcb8d1b65ff6b3622419c9448053c221788e2
Author: Jeff Xu <jeffxu@google.com>
Date:   Wed Jul 5 06:33:15 2023 +0000

    selftests/memfd: sysctl: fix MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED

    [ Upstream commit badbbcd76545c58eff64bb1548f7f834a30dc52a ]

    Add selftest for sysctl vm.memfd_noexec is 2
    (MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED)

    memfd_create(.., MFD_EXEC) should fail in this case.

    Link: https://lkml.kernel.org/r/20230705063315.3680666-3-jeffxu@google.com
    Reported-by: Dominique Martinet <asmadeus@codewreck.org>
    Closes: https://lore.kernel.org/linux-
mm/CABi2SkXUX_QqTQ10Yx9bBUGpN1wByOi_=gZU6WEy5a8MaQY3Jw@mail.gmail.com/T/
    Signed-off-by: Jeff Xu <jeffxu@google.com>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Cc: Hugh Dickins <hughd@google.com>
    Cc: Jann Horn <jannh@google.com>
    Cc: Jorge Lucangeli Obes <jorgelo@chromium.org>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: kernel test robot <lkp@intel.com>
    Cc: Mike Kravetz <mike.kravetz@oracle.com>
    Cc: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Stable-dep-of: 202e14222fad ("memfd: do not -EACCES old memfd_create() users with
vm.memfd_noexec=2")
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 4f417bb14ced2be52c8089b80615767c30cafd77
Author: Jeff Xu <jeffxu@google.com>
Date:   Wed Jul 5 06:33:14 2023 +0000

    mm/memfd: sysctl: fix MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED

    [ Upstream commit 72de259130229412ca49871e70ffaf17dc9fba98 ]

    Patch series "mm/memfd: fix sysctl MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED", v2.

    When sysctl vm.memfd_noexec is 2 (MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED),
    memfd_create(.., MFD_EXEC) should fail.

    This complies with how MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED is defined -
    "memfd_create() without MFD_NOEXEC_SEAL will be rejected"

    Thanks to Dominique Martinet <asmadeus@codewreck.org> who reported the bug.
    see [1] for context.

    [1] https://lore.kernel.org/linux-
mm/CABi2SkXUX_QqTQ10Yx9bBUGpN1wByOi_=gZU6WEy5a8MaQY3Jw@mail.gmail.com/T/

    This patch (of 2):

    When vm.memfd_noexec is 2 (MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED),
    memfd_create(.., MFD_EXEC) should fail.

    This complies with how MEMFD_NOEXEC_SCOPE_NOEXEC_ENFORCED is
    defined - "memfd_create() without MFD_NOEXEC_SEAL will be rejected"

    Link: https://lkml.kernel.org/r/20230705063315.3680666-1-jeffxu@google.com
    Link: https://lkml.kernel.org/r/20230705063315.3680666-2-jeffxu@google.com
    Fixes: 105ff5339f49 ("mm/memfd: add MFD_NOEXEC_SEAL and MFD_EXEC")
    Reported-by: Dominique Martinet <asmadeus@codewreck.org>
    Closes: https://lore.kernel.org/linux-
mm/CABi2SkXUX_QqTQ10Yx9bBUGpN1wByOi_=gZU6WEy5a8MaQY3Jw@mail.gmail.com/T/
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202306301351.kkbSegQW-lkp@intel.com/
    Signed-off-by: Jeff Xu <jeffxu@google.com>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Cc: Hugh Dickins <hughd@google.com>
    Cc: Jann Horn <jannh@google.com>
    Cc: Jorge Lucangeli Obes <jorgelo@chromium.org>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <skhan@linuxfoundation.org>
    Cc: Mike Kravetz <mike.kravetz@oracle.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Stable-dep-of: 202e14222fad ("memfd: do not -EACCES old memfd_create() users with
vm.memfd_noexec=2")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 420be7f23e0cd6c4bc305dd375ed87e35a0b9a5f
Author: Hugo Villeneuve <hvilleneuve@dimonoff.com>
Date:   Mon Aug 7 17:45:54 2023 -0400

    serial: sc16is7xx: fix regression with GPIO configuration

    [ Upstream commit 0499942928341d572a42199580433c2b0725211e ]

    Commit 679875d1d880 ("sc16is7xx: Separate GPIOs from modem control lines")
    and commit 21144bab4f11 ("sc16is7xx: Handle modem status lines")
    changed the function of the GPIOs pins to act as modem control
    lines without any possibility of selecting GPIO function.
```

```
    As a consequence, applications that depends on GPIO lines configured
    by default as GPIO pins no longer work as expected.

    Also, the change to select modem control lines function was done only
    for channel A of dual UART variants (752/762). This was not documented
    in the log message.

    Allow to specify GPIO or modem control line function in the device
    tree, and for each of the ports (A or B).

    Do so by using the new device-tree property named
    "nxp,modem-control-line-ports" (property added in separate patch).

    When registering GPIO chip controller, mask-out GPIO pins declared as
    modem control lines according to this new DT property.

    Fixes: 679875d1d880 ("sc16is7xx: Separate GPIOs from modem control lines")
    Fixes: 21144bab4f11 ("sc16is7xx: Handle modem status lines")
    Cc: stable@vger.kernel.org
    Signed-off-by: Hugo Villeneuve <hvilleneuve@dimonoff.com>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Reviewed-by: Lech Perczak <lech.perczak@camlingroup.com>
    Tested-by: Lech Perczak <lech.perczak@camlingroup.com>
    Acked-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20230807214556.540627-5-hugo@hugovil.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eeeb8a28da00c48a82be82e63101ef362e89a268
Author: Hugo Villeneuve <hvilleneuve@dimonoff.com>
Date:   Mon Aug 7 17:45:52 2023 -0400

    serial: sc16is7xx: remove obsolete out_thread label

    [ Upstream commit dabc54a45711fe77674a6c0348231e00e66bd567 ]

    Commit c8f71b49ee4d ("serial: sc16is7xx: setup GPIO controller later
    in probe") moved GPIO setup code later in probe function. Doing so
    also required to move ports cleanup code (out_ports label) after the
    GPIO cleanup code.

    After these moves, the out_thread label becomes misplaced and makes
    part of the cleanup code illogical.

    This patch remove the now obsolete out_thread label and make GPIO
    setup code jump to out_ports label if it fails.

    Signed-off-by: Hugo Villeneuve <hvilleneuve@dimonoff.com>
    Reviewed-by: Lech Perczak <lech.perczak@camlingroup.com>
    Tested-by: Lech Perczak <lech.perczak@camlingroup.com>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Link: https://lore.kernel.org/r/20230807214556.540627-3-hugo@hugovil.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Stable-dep-of: 049994292834 ("serial: sc16is7xx: fix regression with GPIO
configuration")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 901ce6928ef5533cc5f58450e0d71d4cc4abdebd
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Tue Aug 22 12:02:03 2023 -0700

    Bluetooth: HCI: Introduce HCI_QUIRK_BROKEN_LE_CODED

    [ Upstream commit 253f3399f4c09ce6f4e67350f839be0361b4d5ff ]
```

```
    This introduces HCI_QUIRK_BROKEN_LE_CODED which is used to indicate
    that LE Coded PHY shall not be used, it is then set for some Intel
    models that claim to support it but when used causes many problems.

    Cc: stable@vger.kernel.org # 6.4.y+
    Link: https://github.com/bluez/bluez/issues/577
    Link: https://github.com/bluez/bluez/issues/582
    Link: https://lore.kernel.org/linux-bluetooth/CABBYNZKco-v7wkjHHexxQbgwwSz-
S=GZ=dZKbRE1qxT1h4fFbQ@mail.gmail.com/T/#
    Fixes: 288c90224eec ("Bluetooth: Enable all supported LE PHY by default")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 80fe27811c56ecd3896c3d511ab8298146071d98
Author: Hilda Wu <hildawu@realtek.com>
Date:   Wed Jun 21 18:00:31 2023 +0800

    Bluetooth: msft: Extended monitor tracking by address filter

    [ Upstream commit 9e14606d8f38ea52a38c27692a9c1513c987a5da ]

    Since limited tracking device per condition, this feature is to support
    tracking multiple devices concurrently.
    When a pattern monitor detects the device, this feature issues an address
    monitor for tracking that device. Let pattern monitor can keep monitor
    new devices.
    This feature adds an address filter when receiving a LE monitor device
    event which monitor handle is for a pattern, and the controller started
    monitoring the device. And this feature also has cancelled the monitor
    advertisement from address filters when receiving a LE monitor device
    event when the controller stopped monitoring the device specified by an
    address and monitor handle.

    Below is an example to know the feature adds the address filter.

    //Add MSFT pattern monitor
    < HCI Command: Vendor (0x3f|0x00f0) plen 14          #142 [hci0] 55.552420
            03 b8 a4 03 ff 01 01 06 09 05 5f 52 45 46        ..........._REF
    > HCI Event: Command Complete (0x0e) plen 6          #143 [hci0] 55.653960
          Vendor (0x3f|0x00f0) ncmd 2
            Status: Success (0x00)
            03 00

    //Got event from the pattern monitor
    > HCI Event: Vendor (0xff) plen 18                   #148 [hci0] 58.384953
            23 79 54 33 77 88 97 68 02 00 fb c1 29 eb 27 b8  #yT3w..h....).'.
            00 01                                            ..

    //Add MSFT address monitor (Sample address: B8:27:EB:29:C1:FB)
    < HCI Command: Vendor (0x3f|0x00f0) plen 13          #149 [hci0] 58.385067
            03 b8 a4 03 ff 04 00 fb c1 29 eb 27 b8           .........).'.

    //Report to userspace about found device (ADV Monitor Device Found)
    @ MGMT Event: Unknown (0x002f) plen 38           {0x0003} [hci0] 58.680042
            01 00 fb c1 29 eb 27 b8 01 ce 00 00 00 00 16 00  ....).'.........
            0a 09 4b 45 59 42 44 5f 52 45 46 02 01 06 03 19  ..KEYBD_REF.....
            c1 03 03 03 12 18                                ......

    //Got event from address monitor
    > HCI Event: Vendor (0xff) plen 18                   #152 [hci0] 58.672956
            23 79 54 33 77 88 97 68 02 00 fb c1 29 eb 27 b8  #yT3w..h....).'.
            01 01

    Signed-off-by: Alex Lu <alex_lu@realsil.com.cn>
```

```
    Signed-off-by: Hilda Wu <hildawu@realtek.com>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: 253f3399f4c0 ("Bluetooth: HCI: Introduce HCI_QUIRK_BROKEN_LE_CODED")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4cd71fab288eb03e00f66a20a4297d96ceda056
Author: Arnd Bergmann <arnd@kernel.org>
Date:   Thu Jul 27 14:22:58 2023 +0200

    media: ipu3-cio2: allow ipu_bridge to be a module again

    commit 2545a2c02ba1da9cfb9ec218623c71b00eb4a555 upstream.

    This code was previously part of the VIDEO_IPU3_CIO2 driver, which could
    be built-in or a loadable module, but after the move it turned into a
    builtin-only driver. This fails to link when the I2C subsystem is a
    module:

    x86_64-linux-ld: drivers/media/pci/intel/ipu-bridge.o: in function
`ipu_bridge_unregister_sensors':
    ipu-bridge.c:(.text+0x50): undefined reference to `i2c_unregister_device'
    x86_64-linux-ld: drivers/media/pci/intel/ipu-bridge.o: in function `ipu_bridge_init':
    ipu-bridge.c:(.text+0x9c9): undefined reference to `i2c_acpi_new_device_by_fwnode'

    In general, drivers should not have to be built-in, so change the option
    to a tristate with the corresponding dependency. This in turn opens a
    new problem with the dependency, as the IPU bridge can be a loadable module
    while the ipu3 driver itself is built-in, producing a new link failure:

    86_64-linux-ld: drivers/media/pci/intel/ipu3/ipu3-cio2.o: in function `cio2_pci_probe':
    ipu3-cio2.c:(.text+0x197e): undefined reference to `ipu_bridge_init'

    In order to fix this, restore the old Kconfig option that controlled
    the ipu bridge driver before it was split out, but make it select a
    hidden symbol that now corresponds to the bridge driver.

    When other drivers get added that share ipu-bridge, this should cover
    all corner cases, and allow any combination of them to be built-in
    or modular.

    Link: https://lore.kernel.org/linux-media/20230727122331.2421453-1-arnd@kernel.org

    Fixes: 881ca25978c6 ("media: ipu3-cio2: rename cio2 bridge to ipu bridge and move out of
ipu3")'
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fefce45e4c08888b48066dde2aa327a4b6db6cc6
Author: Kan Liang <kan.liang@linux.intel.com>
Date:   Tue Sep 5 06:42:48 2023 -0700

    perf/x86/uncore: Correct the number of CHAs on EMR

    commit 6f7f984fa85b305799076a1bcec941b9377587de upstream.

    Starting from SPR, the basic uncore PMON information is retrieved from
    the discovery table (resides in an MMIO space populated by BIOS). It is
    called the discovery method. The existing value of the type->num_boxes
    is from the discovery table.

    On some SPR variants, there is a firmware bug that makes the value from the
    discovery table incorrect. We use the value from the
```

```
    SPR_MSR_UNC_CBO_CONFIG MSR to replace the one from the discovery table:

       38776cc45eb7 ("perf/x86/uncore: Correct the number of CHAs on SPR")

    Unfortunately, the SPR_MSR_UNC_CBO_CONFIG isn't available for the EMR
    XCC (Always returns 0), but the above firmware bug doesn't impact the
    EMR XCC.

    Don't let the value from the MSR replace the existing value from the
    discovery table.

    Fixes: 38776cc45eb7 ("perf/x86/uncore: Correct the number of CHAs on SPR")
    Reported-by: Stephane Eranian <eranian@google.com>
    Reported-by: Yunying Sun <yunying.sun@intel.com>
    Signed-off-by: Kan Liang <kan.liang@linux.intel.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Tested-by: Yunying Sun <yunying.sun@intel.com>
    Link: https://lore.kernel.org/r/20230905134248.496114-1-kan.liang@linux.intel.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cc947a043b9da2e221f2abdbb22e5ce291ba4a2d
Author: Song Liu <song@kernel.org>
Date:   Wed Sep 6 10:52:15 2023 -0700

    x86/build: Fix linker fill bytes quirk/incompatibility for ld.lld

    commit 65e710899fd19f435f40268f3a92dfaa11f14470 upstream.

    With ":text =0xcccc", ld.lld fills unused text area with 0xcccc0000.
    Example objdump -D output:

            ffffffff82b04203:       00 00                   add    %al,(%rax)
            ffffffff82b04205:       cc                      int3
            ffffffff82b04206:       cc                      int3
            ffffffff82b04207:       00 00                   add    %al,(%rax)
            ffffffff82b04209:       cc                      int3
            ffffffff82b0420a:       cc                      int3

    Replace it with ":text =0xcccccccc", so we get the following instead:

            ffffffff82b04203:       cc                      int3
            ffffffff82b04204:       cc                      int3
            ffffffff82b04205:       cc                      int3
            ffffffff82b04206:       cc                      int3
            ffffffff82b04207:       cc                      int3
            ffffffff82b04208:       cc                      int3

    gcc/ld doesn't seem to have the same issue. The generated code stays the
    same for gcc/ld.

    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Reviewed-by: Kees Cook <keescook@chromium.org>
    Acked-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Fixes: 7705dc855797 ("x86/vmlinux: Use INT3 instead of NOP for linker fill bytes")
    Link: https://lore.kernel.org/r/20230906175215.2236033-1-song@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bdfa4029f92bc8df6d1f6e59ca3da44a78d9182d
Author: Jack Wang <jinpu.wang@ionos.com>
Date:   Wed Sep 6 15:17:12 2023 +0200

    x86/sgx: Break up long non-preemptible delays in sgx_vepc_release()

    commit 3d7d72a34e05b23e21bafc8bfb861e73c86b31f3 upstream.
```

```
    On large enclaves we hit the softlockup warning with following call trace:

            xa_erase()
            sgx_vepc_release()
            __fput()
            task_work_run()
            do_exit()

    The latency issue is similar to the one fixed in:

      8795359e35bc ("x86/sgx: Silence softlockup detection when releasing large enclaves")

    The test system has 64GB of enclave memory, and all is assigned to a single VM.
    Release of 'vepc' takes a longer time and causes long latencies, which triggers
    the softlockup warning.

    Add cond_resched() to give other tasks a chance to run and reduce
    latencies, which also avoids the softlockup detector.

    [ mingo: Rewrote the changelog. ]

    Fixes: 540745ddbc70 ("x86/sgx: Introduce virtual EPC for use by KVM guests")
    Reported-by: Yu Zhang <yu.zhang@ionos.com>
    Signed-off-by: Jack Wang <jinpu.wang@ionos.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Tested-by: Yu Zhang <yu.zhang@ionos.com>
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Reviewed-by: Kai Huang <kai.huang@intel.com>
    Acked-by: Haitao Huang <haitao.huang@linux.intel.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5674ea8b298b89ff0779a656d5c7df1ec3b74381
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Aug 9 20:52:20 2023 +0200

    x86/smp: Don't send INIT to non-present and non-booted CPUs

    commit 3f874c9b2aae8e30463efc1872bea4baa9ed25dc upstream.

    Vasant reported that kexec() can hang or reset the machine when it tries to
    park CPUs via INIT. This happens when the kernel is using extended APIC,
    but the present mask has APIC IDs >= 0x100 enumerated.

    As extended APIC can only handle 8 bit of APIC ID sending INIT to APIC ID
    0x100 sends INIT to APIC ID 0x0. That's the boot CPU which is special on
    x86 and INIT causes the system to hang or resets the machine.

    Prevent this by sending INIT only to those CPUs which have been booted
    once.

    Fixes: 45e34c8af58f ("x86/smp: Put CPUs into INIT on shutdown if possible")
    Reported-by: Dheeraj Kumar Srivastava <dheerajkumar.srivastava@amd.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Vasant Hegde <vasant.hegde@amd.com>
    Link: https://lore.kernel.org/r/87cyzwjbff.ffs@tglx
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1615e2e021b47302775b9d682c747c85fcf02df2
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 11 13:38:46 2023 -0400

    USB: core: Fix oversight in SuperSpeed initialization
```

```
    commit 59cf445754566984fd55af19ba7146c76e6627bc upstream.

    Commit 85d07c556216 ("USB: core: Unite old scheme and new scheme
    descriptor reads") altered the way USB devices are enumerated
    following detection, and in the process it messed up the
    initialization of SuperSpeed (or faster) devices:

    [   31.650759] usb 2-1: new SuperSpeed Plus Gen 2x1 USB device number 2 using xhci_hcd
    [   31.663107] usb 2-1: device descriptor read/8, error -71
    [   31.952697] usb 2-1: new SuperSpeed Plus Gen 2x1 USB device number 3 using xhci_hcd
    [   31.965122] usb 2-1: device descriptor read/8, error -71
    [   32.080991] usb usb2-port1: attempt power cycle
    ...

    The problem was caused by the commit forgetting that in SuperSpeed or
    faster devices, the device descriptor uses a logarithmic encoding of
    the bMaxPacketSize0 value.  (For some reason I thought the 255 case in
    the switch statement was meant for these devices, but it isn't -- it
    was meant for Wireless USB and is no longer needed.)

    We can fix the oversight by testing for buf->bMaxPacketSize0 = 9
    (meaning 512, the actual maxpacket size for ep0 on all SuperSpeed
    devices) and straightening out the logic that checks and adjusts our
    initial guesses of the maxpacket value.

    Reported-and-tested-by: Thinh Nguyen <Thinh.Nguyen@synopsys.com>
    Closes: https://lore.kernel.org/linux-usb/20230810002257.nadxmfmrobkaxgnz@synopsys.com/
    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Fixes: 85d07c556216 ("USB: core: Unite old scheme and new scheme descriptor reads")
    Link: https://lore.kernel.org/r/8809e6c5-59d5-4d2d-ac8f-6d106658ad73@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 43fe560165353646830eb95349503e554e3eeb34
Author: Douglas Anderson <dianders@chromium.org>
Date:   Thu Jul 27 10:16:31 2023 -0700

    of: property: fw_devlink: Add a devlink for panel followers

    commit fbf0ea2da3c7cd0b33ed7ae53a67ab1c24838cba upstream.

    Inform fw_devlink of the fact that a panel follower (like a
    touchscreen) is effectively a consumer of the panel from the purposes
    of fw_devlink.

    NOTE: this patch isn't required for correctness but instead optimizes
    probe order / helps avoid deferrals.

    Acked-by: Rob Herring <robh@kernel.org>
    Reviewed-by: Maxime Ripard <mripard@kernel.org>
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Link:
https://patchwork.freedesktop.org/patch/msgid/20230727101636.v4.4.Ibf8e1342b5b7906279db2365a
ca45e6253857bb3@changeid
    Cc: Adam Ford <aford173@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 98e8f162350558c624ca1234d569ffa712ea3495
Author: Gustavo A. R. Silva <gustavoars@kernel.org>
Date:   Mon Jul 31 21:15:48 2023 -0600

    cpufreq: brcmstb-avs-cpufreq: Fix -Warray-bounds bug

    commit e520d0b6be950ce3738cf4b9bd3b392be818f1dc upstream.

    Allocate extra space for terminating element at:
```

```
    drivers/cpufreq/brcmstb-avs-cpufreq.c:
    449         table[i].frequency = CPUFREQ_TABLE_END;

    and add code comment to make this clear.

    This fixes the following -Warray-bounds warning seen after building
    ARM with multi_v7_defconfig (GCC 13):
    In function 'brcm_avs_get_freq_table',
        inlined from 'brcm_avs_cpufreq_init' at drivers/cpufreq/brcmstb-avs-
cpufreq.c:623:15:
    drivers/cpufreq/brcmstb-avs-cpufreq.c:449:28: warning: array subscript 5 is outside
array bounds of 'void[60]' [-Warray-bounds=]
      449 |         table[i].frequency = CPUFREQ_TABLE_END;
    In file included from include/linux/node.h:18,
                     from include/linux/cpu.h:17,
                     from include/linux/cpufreq.h:12,
                     from drivers/cpufreq/brcmstb-avs-cpufreq.c:44:
    In function 'devm_kmalloc_array',
        inlined from 'devm_kcalloc' at include/linux/device.h:328:9,
        inlined from 'brcm_avs_get_freq_table' at drivers/cpufreq/brcmstb-avs-
cpufreq.c:437:10,
        inlined from 'brcm_avs_cpufreq_init' at drivers/cpufreq/brcmstb-avs-
cpufreq.c:623:15:
    include/linux/device.h:323:16: note: at offset 60 into object of size 60 allocated by
'devm_kmalloc'
      323 |         return devm_kmalloc(dev, bytes, flags);
          |                ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    This helps with the ongoing efforts to tighten the FORTIFY_SOURCE
    routines on memcpy() and help us make progress towards globally
    enabling -Warray-bounds.

    Link: https://github.com/KSPP/linux/issues/324
    Fixes: de322e085995 ("cpufreq: brcmstb-avs-cpufreq: AVS CPUfreq driver for Broadcom STB
SoCs")
    Cc: stable@vger.kernel.org
    Signed-off-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Reviewed-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 11002a62f07ad33c79d96352d99be0e63aca2064
Author: Thomas Bourgoin <thomas.bourgoin@foss.st.com>
Date:   Thu Jul 13 17:15:17 2023 +0200

    crypto: stm32 - fix MDMAT condition

    commit a4adfbc2544933ac12e7fbd50708290265546dbc upstream.

    If IP has MDMAT support, set or reset the bit MDMAT in Control Register.

    Fixes: b56403a25af7 ("crypto: stm32/hash - Support Ux500 hash")
    Cc: stable@vger.kernel.org
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Thomas Bourgoin <thomas.bourgoin@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 861c82cda5f12156525366ede98e106ce6b976b9
Author: Thomas Bourgoin <thomas.bourgoin@foss.st.com>
Date:   Thu Jul 13 17:15:15 2023 +0200

    crypto: stm32 - fix loop iterating through scatterlist for DMA
```

```
    commit d9c83f71eeceed2cb54bb78be84f2d4055fd9a1f upstream.

    We were reading the length of the scatterlist sg after copying value of
    tsg inside.
    So we are using the size of the previous scatterlist and for the first
    one we are using an unitialised value.
    Fix this by copying tsg in sg[0] before reading the size.

    Fixes : 8a1012d3f2ab ("crypto: stm32 - Support for STM32 HASH module")
    Cc: stable@vger.kernel.org
    Signed-off-by: Thomas Bourgoin <thomas.bourgoin@foss.st.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f617ab148203eb68145ff6cff22903b6696e1d35
Author: Benjamin Tissoires <benjamin.tissoires@redhat.com>
Date:   Wed Jul 12 17:02:34 2023 +0200

    HID: logitech-hidpp: rework one more time the retries attempts

    commit 60165ab774cb0c509680a73cf826d0e158454653 upstream.

    Extract the internal code inside a helper function, fix the
    initialization of the parameters used in the helper function
    (`hidpp->answer_available` was not reset and `*response` wasn't either),
    and use a `do {...} while();` loop.

    Fixes: 586e8fede795 ("HID: logitech-hidpp: Retry commands when device is busy")
    Cc: stable@vger.kernel.org
    Reviewed-by: Bastien Nocera <hadess@hadess.net>
    Signed-off-by: Benjamin Tissoires <benjamin.tissoires@redhat.com>
    Link: https://lore.kernel.org/r/20230621-logitech-fixes-v2-1-3635f7f9c8af@kernel.org
    Signed-off-by: Benjamin Tissoires <bentiss@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 614cc44384bf4692e2d545ee9c830fcfc991e11f
Author: Heiko Carstens <hca@linux.ibm.com>
Date:   Mon Aug 28 17:31:42 2023 +0200

    s390/dasd: fix string length handling

    commit f7cf22424665043787a96a66a048ff6b2cfd473c upstream.

    Building dasd_eckd.o with latest clang reveals this bug:

        CC      drivers/s390/block/dasd_eckd.o
          drivers/s390/block/dasd_eckd.c:1082:3: warning: 'snprintf' will always be
truncated;
          specified size is 1, but format string expands to at least 11 [-Wfortify-source]
           1082 |                 snprintf(print_uid, sizeof(*print_uid),
                |                 ^
          drivers/s390/block/dasd_eckd.c:1087:3: warning: 'snprintf' will always be
truncated;
          specified size is 1, but format string expands to at least 10 [-Wfortify-source]
           1087 |                 snprintf(print_uid, sizeof(*print_uid),
                |                 ^

    Fix this by moving and using the existing UID_STRLEN for the arrays
    that are being written to. Also rename UID_STRLEN to DASD_UID_STRLEN
    to clarify its scope.

    Fixes: 23596961b437 ("s390/dasd: split up dasd_eckd_read_conf")
    Reviewed-by: Peter Oberparleiter <oberpar@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Tested-by: Nick Desaulniers <ndesaulniers@google.com> # build
```

```
    Reported-by: Nathan Chancellor <nathan@kernel.org>
    Closes: https://github.com/ClangBuiltLinux/linux/issues/1923
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Link: https://lore.kernel.org/r/20230828153142.2843753-2-hca@linux.ibm.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e938a5aec082b9978ec202bbe8f0c4a3f2a524f5
Author: Sven Schnelle <svens@linux.ibm.com>
Date:   Tue Aug 15 09:26:06 2023 +0200

    s390/ipl: add missing secure/has_secure file to ipl type 'unknown'

    commit ea5717cb13468323a7c3dd394748301802991f39 upstream.

    OS installers are relying on /sys/firmware/ipl/has_secure to be
    present on machines supporting secure boot. This file is present
    for all IPL types, but not the unknown type, which prevents a secure
    installation when an LPAR is booted in HMC via FTP(s), because
    this is an unknown IPL type in linux. While at it, also add the secure
    file.

    Fixes: c9896acc7851 ("s390/ipl: Provide has_secure sysfs attribute")
    Cc: stable@vger.kernel.org
    Signed-off-by: Sven Schnelle <svens@linux.ibm.com>
    Reviewed-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b5c531a9a7d8e047c90c909f09cef06a9f8e62f4
Author: Gerald Schaefer <gerald.schaefer@linux.ibm.com>
Date:   Thu Aug 10 10:22:36 2023 +0200

    s390/dcssblk: fix kernel crash with list_add corruption

    commit c8f40a0bccefd613748d080147469a4652d6e74c upstream.

    Commit fb08a1908cb1 ("dax: simplify the dax_device <-> gendisk
    association") introduced new logic for gendisk association, requiring
    drivers to explicitly call dax_add_host() and dax_remove_host().

    For dcssblk driver, some dax_remove_host() calls were missing, e.g. in
    device remove path. The commit also broke error handling for out_dax case
    in device add path, resulting in an extra put_device() w/o the previous
    get_device() in that case.

    This lead to stale xarray entries after device add / remove cycles. In the
    case when a previously used struct gendisk pointer (xarray index) would be
    used again, because blk_alloc_disk() happened to return such a pointer, the
    xa_insert() in dax_add_host() would fail and go to out_dax, doing the extra
    put_device() in the error path. In combination with an already flawed error
    handling in dcssblk (device_register() cleanup), which needs to be
    addressed in a separate patch, this resulted in a missing device_del() /
    klist_del(), and eventually in the kernel crash with list_add corruption on
    a subsequent device_add() / klist_add().

    Fix this by adding the missing dax_remove_host() calls, and also move the
    put_device() in the error path to restore the previous logic.

    Fixes: fb08a1908cb1 ("dax: simplify the dax_device <-> gendisk association")
    Cc: <stable@vger.kernel.org> # 5.17+
    Acked-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Gerald Schaefer <gerald.schaefer@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```

```
commit f2a9cd050901c4734bc175294c3bd6a40a1a9e99
Author: Andy Chiu <andy.chiu@sifive.com>
Date:   Fri Aug 25 05:02:46 2023 +0000

    RISC-V: Add ptrace support for vectors

    commit 9300f00439743c4a34d735e1a27118eb68a1504e upstream.

    This patch add back the ptrace support with the following fix:
     - Define NT_RISCV_CSR and re-number NT_RISCV_VECTOR to prevent
       conflicting with gdb's NT_RISCV_CSR.
     - Use struct __riscv_v_regset_state to handle ptrace requests

    Since gdb does not directly include the note description header in
    Linux and has already defined NT_RISCV_CSR as 0x900, we decide to
    sync with gdb and renumber NT_RISCV_VECTOR to solve and prevent future
    conflicts.

    Fixes: 0c59922c769a ("riscv: Add ptrace vector support")
    Signed-off-by: Andy Chiu <andy.chiu@sifive.com>
    Link: https://lore.kernel.org/r/20230825050248.32681-1-andy.chiu@sifive.com
    [Palmer: Drop the unused "size" variable in riscv_vr_set().]
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 69a7ff3001fcb78f7defcaefddc47aa6841cf1f2
Author: David Howells <dhowells@redhat.com>
Date:   Fri Sep 8 17:03:20 2023 +0100

    iov_iter: Fix iov_iter_extract_pages() with zero-sized entries

    commit f741bd7178c95abd7aeac5a9d933ee542f9a5509 upstream.

    iov_iter_extract_pages() doesn't correctly handle skipping over initial
    zero-length entries in ITER_KVEC and ITER_BVEC-type iterators.

    The problem is that it accidentally reduces maxsize to 0 when it
    skipping and thus runs to the end of the array and returns 0.

    Fix this by sticking the calculated size-to-copy in a new variable
    rather than back in maxsize.

    Fixes: 7d58fe731028 ("iov_iter: Add a function to extract a page list from an iterator")
    Signed-off-by: David Howells <dhowells@redhat.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Jens Axboe <axboe@kernel.dk>
    Cc: Al Viro <viro@zeniv.linux.org.uk>
    Cc: David Hildenbrand <david@redhat.com>
    Cc: John Hubbard <jhubbard@nvidia.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d8a743b0b0b14f50a415d911d2a70c940b2a4e98
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Tue Jul 25 18:40:47 2023 +0200

    regulator: dt-bindings: qcom,rpm: fix pattern for children

    commit 75d9bf03e2fa38242b35e941ce7c7cdabe479961 upstream.

    The "or" (|) in regular expression must be within parentheses,
    otherwise it is not really an "or" and it matches supplies:
```

```
       qcom-apq8060-dragonboard.dtb: regulators-1: vdd_ncp-supply: [[34]] is not of type
  'object'

      Fixes: fde0e25b71a9 ("dt-bindings: regulators: convert non-smd RPM Regulators bindings
  to dt-schema")
      Cc: stable@vger.kernel.org
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
      Link: https://lore.kernel.org/r/20230725164047.368892-1-krzysztof.kozlowski@linaro.org
      Signed-off-by: Mark Brown <broonie@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 4ed3664fbc7ab23fed55b6a1f2dd6af71ff7e51d
  Author: D Scott Phillips <scott@os.amperecomputing.com>
  Date:   Mon Jun 26 17:29:39 2023 -0700

      arm64: sdei: abort running SDEI handlers during crash

      commit 5cd474e57368f0957c343bb21e309cf82826b1ef upstream.

      Interrupts are blocked in SDEI context, per the SDEI spec: "The client
      interrupts cannot preempt the event handler." If we crashed in the SDEI
      handler-running context (as with ACPI's AGDI) then we need to clean up the
      SDEI state before proceeding to the crash kernel so that the crash kernel
      can have working interrupts.

      Track the active SDEI handler per-cpu so that we can COMPLETE_AND_RESUME
      the handler, discarding the interrupted context.

      Fixes: f5df26961853 ("arm64: kernel: Add arch-specific SDEI entry code and CPU masking")
      Signed-off-by: D Scott Phillips <scott@os.amperecomputing.com>
      Cc: stable@vger.kernel.org
      Reviewed-by: James Morse <james.morse@arm.com>
      Tested-by: Mihai Carabas <mihai.carabas@oracle.com>
      Link: https://lore.kernel.org/r/20230627002939.2758-1-scott@os.amperecomputing.com
      Signed-off-by: Will Deacon <will@kernel.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 25fb4e3402d46f425ec135ef6f09792a4c1b3003
  Author: Enlin Mu <enlin.mu@unisoc.com>
  Date:   Tue Aug 1 14:04:32 2023 +0800

      pstore/ram: Check start of empty przs during init

      commit fe8c3623ab06603eb760444a032d426542212021 upstream.

      After commit 30696378f68a ("pstore/ram: Do not treat empty buffers as
      valid"), initialization would assume a prz was valid after seeing that
      the buffer_size is zero (regardless of the buffer start position). This
      unchecked start value means it could be outside the bounds of the buffer,
      leading to future access panics when written to:

       sysdump_panic_event+0x3b4/0x5b8
       atomic_notifier_call_chain+0x54/0x90
       panic+0x1c8/0x42c
       die+0x29c/0x2a8
       die_kernel_fault+0x68/0x78
       __do_kernel_fault+0x1c4/0x1e0
       do_bad_area+0x40/0x100
       do_translation_fault+0x68/0x80
       do_mem_abort+0x68/0xf8
       el1_da+0x1c/0xc0
       __raw_writeb+0x38/0x174
       __memcpy_toio+0x40/0xac
       persistent_ram_update+0x44/0x12c
       persistent_ram_write+0x1a8/0x1b8
```

```
       ramoops_pstore_write+0x198/0x1e8
       pstore_console_write+0x94/0xe0
       ...

    To avoid this, also check if the prz start is 0 during the initialization
    phase. If not, the next prz sanity check case will discover it (start >
    size) and zap the buffer back to a sane state.

    Fixes: 30696378f68a ("pstore/ram: Do not treat empty buffers as valid")
    Cc: Yunlong Xing <yunlong.xing@unisoc.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Enlin Mu <enlin.mu@unisoc.com>
    Link: https://lore.kernel.org/r/20230801060432.1307717-1-yunlong.xing@unisoc.com
    [kees: update commit log with backtrace and clarifications]
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit afd3199395899ef51adf3758c44d0ca8d648f38f
Author: Wolfram Sang <wsa+renesas@sang-engineering.com>
Date:   Wed Jul 12 16:00:11 2023 +0200

    mmc: renesas_sdhi: register irqs before registering controller

    commit 74f45de394d979cc7770271f92fafa53e1ed3119 upstream.

    IRQs should be ready to serve when we call mmc_add_host() via
    tmio_mmc_host_probe(). To achieve that, ensure that all irqs are masked
    before registering the handlers.

    Signed-off-by: Wolfram Sang <wsa+renesas@sang-engineering.com>
    Tested-by: Biju Das <biju.das.jz@bp.renesas.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Tested-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230712140011.18602-1-wsa+renesas@sang-engineering.com
    Signed-off-by: Ulf Hansson <ulf.hansson@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 609a37db9838c3b8f0be6aea13b985107c6a0f51
Author: Tzung-Bi Shih <tzungbi@kernel.org>
Date:   Thu Aug 3 09:12:45 2023 +0800

    platform/chrome: chromeos_acpi: print hex string for ACPI_TYPE_BUFFER

    commit 0820debb7d489e9eb1f68b7bb69e6ae210699b3f upstream.

    `element->buffer.pointer` should be binary blob.  `%s` doesn't work
    perfect for them.

    Print hex string for ACPI_TYPE_BUFFER.  Also update the documentation
    to reflect this.

    Fixes: 0a4cad9c11ad ("platform/chrome: Add ChromeOS ACPI device driver")
    Cc: stable@vger.kernel.org
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20230803011245.3773756-1-tzungbi@kernel.org
    Signed-off-by: Tzung-Bi Shih <tzungbi@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a424ccd65316502717ab18e9a5bd483bcc45716d
Author: Frederick Lawler <fred@cloudflare.com>
Date:   Tue Aug 1 08:57:09 2023 -0500

    crypto: af_alg - Decrement struct key.usage in alg_set_by_key_serial()
```

```
    commit 6b4b53ca0b7300ba2af98a49dbce22054bf034fe upstream.

    Calls to lookup_user_key() require a corresponding key_put() to
    decrement the usage counter. Once it reaches zero, we schedule key GC.
    Therefore decrement struct key.usage in alg_set_by_key_serial().

    Fixes: 7984ceb134bf ("crypto: af_alg - Support symmetric encryption via keyring keys")
    Cc: <stable@vger.kernel.org>
    Signed-off-by: Frederick Lawler <fred@cloudflare.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2e01bdf7203c383e9d8489d9f963c52d6c81e4db
Author: Yazen Ghannam <yazen.ghannam@amd.com>
Date:   Mon Aug 14 15:08:53 2023 -0500

    x86/MCE: Always save CS register on AMD Zen IF Poison errors

    commit 4240e2ebe67941ce2c4f5c866c3af4b5ac7a0c67 upstream.

    The Instruction Fetch (IF) units on current AMD Zen-based systems do not
    guarantee a synchronous #MC is delivered for poison consumption errors.
    Therefore, MCG_STATUS[EIPV|RIPV] will not be set. However, the
    microarchitecture does guarantee that the exception is delivered within
    the same context. In other words, the exact rIP is not known, but the
    context is known to not have changed.

    There is no architecturally-defined method to determine this behavior.

    The Code Segment (CS) register is always valid on such IF unit poison
    errors regardless of the value of MCG_STATUS[EIPV|RIPV].

    Add a quirk to save the CS register for poison consumption from the IF
    unit banks.

    This is needed to properly determine the context of the error.
    Otherwise, the severity grading function will assume the context is
    IN_KERNEL due to the m->cs value being 0 (the initialized value). This
    leads to unnecessary kernel panics on data poison errors due to the
    kernel believing the poison consumption occurred in kernel context.

    Signed-off-by: Yazen Ghannam <yazen.ghannam@amd.com>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230814200853.29258-1-yazen.ghannam@amd.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5b04dbb97f1090776ca668be77ea4fd635f1531b
Author: Eric Biggers <ebiggers@google.com>
Date:   Tue Aug 1 21:03:53 2023 -0700

    fsverity: skip PKCS#7 parser when keyring is empty

    commit 919dc320956ea353a7fb2d84265195ad5ef525ac upstream.

    If an fsverity builtin signature is given for a file but the
    ".fs-verity" keyring is empty, there's no real reason to run the PKCS#7
    parser.  Skip this to avoid the PKCS#7 attack surface when builtin
    signature support is configured into the kernel but is not being used.

    This is a hardening improvement, not a fix per se, but I've added
    Fixes and Cc stable to get it out to more users.

    Fixes: 432434c9f8e1 ("fs-verity: support builtin file signatures")
    Cc: stable@vger.kernel.org
```

```
    Reviewed-by: Jarkko Sakkinen <jarkko@kernel.org>
    Link: https://lore.kernel.org/r/20230820173237.2579-1-ebiggers@kernel.org
    Signed-off-by: Eric Biggers <ebiggers@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d68331f61768488941c792f13b00b371da3ab2b6
Author: Nicolas Dichtel <nicolas.dichtel@6wind.com>
Date:   Wed Aug 23 15:41:02 2023 +0200

    net: handle ARPHRD_PPP in dev_is_mac_header_xmit()

    commit a4f39c9f14a634e4cd35fcd338c239d11fcc73fc upstream.

    The goal is to support a bpf_redirect() from an ethernet device (ingress)
    to a ppp device (egress).
    The l2 header is added automatically by the ppp driver, thus the ethernet
    header should be removed.

    CC: stable@vger.kernel.org
    Fixes: 27b29f63058d ("bpf: add bpf_redirect() helper")
    Signed-off-by: Nicolas Dichtel <nicolas.dichtel@6wind.com>
    Tested-by: Siwar Zitouni <siwar.zitouni@6wind.com>
    Reviewed-by: Guillaume Nault <gnault@redhat.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e73a3c788735fcaba7f812af9b91409137a27d1a
Author: Thore Sommer <public@thson.de>
Date:   Tue Aug 15 14:29:42 2023 +0300

    X.509: if signature is unsupported skip validation

    commit ef5b52a631f8c18353e80ccab8408b963305510c upstream.

    When the hash algorithm for the signature is not available the digest size
    is 0 and the signature in the certificate is marked as unsupported.

    When validating a self-signed certificate, this needs to be checked,
    because otherwise trying to validate the signature will fail with an
    warning:

    Loading compiled-in X.509 certificates
    WARNING: CPU: 0 PID: 1 at crypto/rsa-pkcs1pad.c:537 \
    pkcs1pad_verify+0x46/0x12c
    ...
    Problem loading in-kernel X.509 certificate (-22)

    Signed-off-by: Thore Sommer <public@thson.de>
    Cc: stable@vger.kernel.org # v4.7+
    Fixes: 6c2dc5ae4ab7 ("X.509: Extract signature digest and make self-signed cert checks
earlier")
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 178619e0b55b4c1ba3c7ad3baa1acf3e8bd6906b
Author: Heiner Kallweit <hkallweit1@gmail.com>
Date:   Fri Aug 25 21:44:01 2023 +0200

    r8169: fix ASPM-related issues on a number of systems with NIC version from RTL8168h

    commit 90ca51e8c654699b672ba61aeaa418dfb3252e5e upstream.

    This effectively reverts 4b5f82f6aaef. On a number of systems ASPM L1
    causes tx timeouts with RTL8168h, see referenced bug report.
```

```
    Fixes: 4b5f82f6aaef ("r8169: enable ASPM L1/L1.1 from RTL8168h")
    Cc: stable@vger.kernel.org
    Closes: https://bugzilla.kernel.org/show_bug.cgi?id=217814
    Signed-off-by: Heiner Kallweit <hkallweit1@gmail.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8ae7457e71a320867d868f2622d7c643596e4f43
Author: Steve Rutherford <srutherford@google.com>
Date:   Thu Aug 24 15:37:31 2023 -0700

    x86/sev: Make enc_dec_hypercall() accept a size instead of npages

    commit ac3f9c9f1b37edaa7d1a9b908bc79d843955a1a2 upstream.

    enc_dec_hypercall() accepted a page count instead of a size, which
    forced its callers to round up. As a result, non-page aligned
    vaddrs caused pages to be spuriously marked as decrypted via the
    encryption status hypercall, which in turn caused consistent
    corruption of pages during live migration. Live migration requires
    accurate encryption status information to avoid migrating pages
    from the wrong perspective.

    Fixes: 064ce6c550a0 ("mm: x86: Invoke hypercall when page encryption status is changed")
    Signed-off-by: Steve Rutherford <srutherford@google.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Reviewed-by: Tom Lendacky <thomas.lendacky@amd.com>
    Reviewed-by: Pankaj Gupta <pankaj.gupta@amd.com>
    Tested-by: Ben Hillier <bhillier@google.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230824223731.2055016-1-srutherford@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec620c34f5fa5d055f9f6136a387755db6157712
Author: Jann Horn <jannh@google.com>
Date:   Fri Aug 25 15:32:41 2023 +0200

    dccp: Fix out of bounds access in DCCP error handler

    commit 977ad86c2a1bcaf58f01ab98df5cc145083c489c upstream.

    There was a previous attempt to fix an out-of-bounds access in the DCCP
    error handlers, but that fix assumed that the error handlers only want
    to access the first 8 bytes of the DCCP header. Actually, they also look
    at the DCCP sequence number, which is stored beyond 8 bytes, so an
    explicit pskb_may_pull() is required.

    Fixes: 6706a97fec96 ("dccp: fix out of bound access in dccp_v4_err()")
    Fixes: 1aa9d1a0e7ee ("ipv6: dccp: fix out of bound access in dccp_v6_err()")
    Cc: stable@vger.kernel.org
    Signed-off-by: Jann Horn <jannh@google.com>
    Reviewed-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 999436d17b4928ca21f67a86afa5f0fe6e816d8a
Author: Alexander Aring <aahringo@redhat.com>
Date:   Thu Aug 24 16:51:42 2023 -0400

    dlm: fix plock lookup when using multiple lockspaces

    commit 7c53e847ff5e97f033fdd31f71949807633d506b upstream.

    All posix lock ops, for all lockspaces (gfs2 file systems) are
    sent to userspace (dlm_controld) through a single misc device.
```

```
    The dlm_controld daemon reads the ops from the misc device
    and sends them to other cluster nodes using separate, per-lockspace
    cluster api communication channels.  The ops for a single lockspace
    are ordered at this level, so that the results are received in
    the same sequence that the requests were sent.  When the results
    are sent back to the kernel via the misc device, they are again
    funneled through the single misc device for all lockspaces.  When
    the dlm code in the kernel processes the results from the misc
    device, these results will be returned in the same sequence that
    the requests were sent, on a per-lockspace basis.  A recent change
    in this request/reply matching code missed the "per-lockspace"
    check (fsid comparison) when matching request and reply, so replies
    could be incorrectly matched to requests from other lockspaces.

    Cc: stable@vger.kernel.org
    Reported-by: Barry Marson <bmarson@redhat.com>
    Fixes: 57e2c2f2d94c ("fs: dlm: fix mismatch of plock results from userspace")
    Signed-off-by: Alexander Aring <aahringo@redhat.com>
    Signed-off-by: David Teigland <teigland@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit acfdc8b77016c8e648aadc283177546c88083dd3
Author: Yafang Shao <laoar.shao@gmail.com>
Date:   Wed Aug 23 02:07:02 2023 +0000

    bpf: Fix issue in verifying allow_ptr_leaks

    commit d75e30dddf73449bc2d10bb8e2f1a2c446bc67a2 upstream.

    After we converted the capabilities of our networking-bpf program from
    cap_sys_admin to cap_net_admin+cap_bpf, our networking-bpf program
    failed to start. Because it failed the bpf verifier, and the error log
    is "R3 pointer comparison prohibited".

    A simple reproducer as follows,

    SEC("cls-ingress")
    int ingress(struct __sk_buff *skb)
    {
            struct iphdr *iph = (void *)(long)skb->data + sizeof(struct ethhdr);

            if ((long)(iph + 1) > (long)skb->data_end)
                    return TC_ACT_STOLEN;
            return TC_ACT_OK;
    }

    Per discussion with Yonghong and Alexei [1], comparison of two packet
    pointers is not a pointer leak. This patch fixes it.

    Our local kernel is 6.1.y and we expect this fix to be backported to
    6.1.y, so stable is CCed.

    [1]. https://lore.kernel.org/bpf/CAADnVQ+Nmspr7Si+pxWn8zkE7hX-
7s93ugwC+94aXSy4uQ9vBg@mail.gmail.com/

    Suggested-by: Yonghong Song <yonghong.song@linux.dev>
    Suggested-by: Alexei Starovoitov <alexei.starovoitov@gmail.com>
    Signed-off-by: Yafang Shao <laoar.shao@gmail.com>
    Acked-by: Eduard Zingerman <eddyz87@gmail.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230823020703.3790-2-laoar.shao@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a3199a11935c2d2a9c85670e1c241c5c46cba02b
```

```
Author: Fudong Wang <fudong.wang@amd.com>
Date:   Fri Aug 11 08:24:59 2023 +0800

    drm/amd/display: Add smu write msg id fail retry process

    commit 72105dcfa3d12b5af49311f857e3490baa225135 upstream.

    A benchmark stress test (12-40 machines x 48hours) found that DCN315 has
    cases where DC writes to an indirect register to set the smu clock msg
    id, but when we go to read the same indirect register the returned msg
    id doesn't match with what we just set it to. So, to fix this retry the
    write until the register's value matches with the requested value.

    Cc: stable@vger.kernel.org # 6.1+
    Fixes: f94903996140 ("drm/amd/display: Add DCN315 CLK_MGR")
    Reviewed-by: Charlene Liu <charlene.liu@amd.com>
    Acked-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Fudong Wang <fudong.wang@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8fce2f8e5b432b7bc5d040a792f61c8382b00dc7
Author: Ekansh Gupta <quic_ekangupt@quicinc.com>
Date:   Fri Aug 11 12:56:43 2023 +0100

    misc: fastrpc: Pass proper scm arguments for static process init

    commit fe6518d547fc52ba74201018dc9aeb364072ac78 upstream.

    Memory is allocated for dynamic loading when audio daemon is trying
    to attach to audioPD on DSP side. This memory is allocated from
    reserved CMA memory region and needs ownership assignment to
    new VMID in order to use it from audioPD.

    In the current implementation, arguments are not correctly passed
    to the scm call which might result in failure of dynamic loading
    on audioPD. Added changes to pass correct arguments during daemon
    attach request.

    Fixes: 0871561055e6 ("misc: fastrpc: Add support for audiopd")
    Cc: stable <stable@kernel.org>
    Tested-by: Ekansh Gupta <quic_ekangupt@quicinc.com>
    Signed-off-by: Ekansh Gupta <quic_ekangupt@quicinc.com>
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230811115643.38578-4-srinivas.kandagatla@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 2b1612ea29959c12cf2bd1b5554ab08759537051
Author: Helge Deller <deller@gmx.de>
Date:   Fri Aug 18 22:48:04 2023 +0200

    parisc: Fix /proc/cpuinfo output for lscpu

    commit 9f5ba4b3e1b3c123eeca5d2d09161e8720048b5c upstream.

    The lscpu command is broken since commit cab56b51ec0e ("parisc: Fix
    device names in /proc/iomem") added the PA pathname to all PA
    devices, includig the CPUs.

    lscpu parses /proc/cpuinfo and now believes it found different CPU
    types since every CPU is listed with an unique identifier (PA
    pathname).

    Fix this problem by simply dropping the PA pathname when listing the
    CPUs in /proc/cpuinfo. There is no need to show the pathname in this
```

```
    procfs file.

    Fixes: cab56b51ec0e ("parisc: Fix device names in /proc/iomem")
    Signed-off-by: Helge Deller <deller@gmx.de>
    Cc: <stable@vger.kernel.org> # v4.9+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a08eaac6652b4a9e8708f6840c8076bd5b234b4c
Author: Aleksa Sarai <cyphar@cyphar.com>
Date:   Fri Jul 14 00:09:58 2023 +1000

    procfs: block chmod on /proc/thread-self/comm

    commit ccf61486fe1e1a48e18c638d1813cda77b3c0737 upstream.

    Due to an oversight in commit 1b3044e39a89 ("procfs: fix pthread
    cross-thread naming if !PR_DUMPABLE") in switching from REG to NOD,
    chmod operations on /proc/thread-self/comm were no longer blocked as
    they are on almost all other procfs files.

    A very similar situation with /proc/self/environ was used to as a root
    exploit a long time ago, but procfs has SB_I_NOEXEC so this is simply a
    correctness issue.

    Ref: https://lwn.net/Articles/191954/
    Ref: 6d76fa58b050 ("Don't allow chmod() on the /proc/<pid>/ files")
    Fixes: 1b3044e39a89 ("procfs: fix pthread cross-thread naming if !PR_DUMPABLE")
    Cc: stable@vger.kernel.org # v4.7+
    Signed-off-by: Aleksa Sarai <cyphar@cyphar.com>
    Message-Id: <20230713141001.27046-1-cyphar@cyphar.com>
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f42cee0051db25ab6d0da4348f00729f22b5493f
Author: Li Lingfeng <lilingfeng3@huawei.com>
Date:   Thu Aug 31 15:59:00 2023 +0800

    block: don't add or resize partition on the disk with GENHD_FL_NO_PART

    commit 1a721de8489fa559ff4471f73c58bb74ac5580d3 upstream.

    Commit a33df75c6328 ("block: use an xarray for disk->part_tbl") remove
    disk_expand_part_tbl() in add_partition(), which means all kinds of
    devices will support extended dynamic `dev_t`.
    However, some devices with GENHD_FL_NO_PART are not expected to add or
    resize partition.
    Fix this by adding check of GENHD_FL_NO_PART before add or resize
    partition.

    Fixes: a33df75c6328 ("block: use an xarray for disk->part_tbl")
    Signed-off-by: Li Lingfeng <lilingfeng3@huawei.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230831075900.1725842-1-lilingfeng@huaweicloud.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ccd700dccea23fb837470902466d52b010eed19c
Author: Christoph Hellwig <hch@lst.de>
Date:   Tue Sep 5 14:47:31 2023 +0200

    block: fix pin count management when merging same-page segments

    commit 5905afc2c7bb713d52c7c7585565feecbb686b44 upstream.

    There is no need to unpin the added page when adding it to the bio fails
```

```
    as that is done by the loop below.  Instead we want to unpin it when adding
    a single page to the bio more than once as bio_release_pages will only
    unpin it once.

    Fixes: d1916c86ccdc ("block: move same page handling from __bio_add_pc_page to the
callers")
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Damien Le Moal <dlemoal@kernel.org>
    Link: https://lore.kernel.org/r/20230905124731.328255-1-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 186958905067391ec123170980a94a653d072b8d
Author: Bjorn Helgaas <bhelgaas@google.com>
Date:   Fri Sep 8 14:55:30 2023 -0500

    Revert "PCI: Mark NVIDIA T4 GPUs to avoid bus reset"

    commit 5260bd6d36c83c5b269c33baaaf8c78e520908b0 upstream.

    This reverts commit d5af729dc2071273f14cbb94abbc60608142fd83.

    d5af729dc207 ("PCI: Mark NVIDIA T4 GPUs to avoid bus reset") avoided
    Secondary Bus Reset on the T4 because the reset seemed to not work when the
    T4 was directly attached to a Root Port.

    But NVIDIA thinks the issue is probably related to some issue with the Root
    Port, not with the T4.  The T4 provides neither PM nor FLR reset, so
    masking bus reset compromises this device for assignment scenarios.

    Revert d5af729dc207 as requested by Wu Zongyong.  This will leave SBR
    broken in the specific configuration Wu tested, as it was in v6.5, so Wu
    will debug that further.

    Link: https://lore.kernel.org/r/ZPqMCDWvITlOLHgJ@wuzongyong-alibaba
    Link: https://lore.kernel.org/r/20230908201104.GA305023@bhelgaas
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b913f4ffa2f7447eed836b4a5252a12ba755eda8
Author: Dave Jiang <dave.jiang@intel.com>
Date:   Tue Aug 22 09:04:57 2023 -0700

    ntb: Fix calculation ntb_transport_tx_free_entry()

    commit 5a7693e6bbf19b22fd6c1d2c4b7beb0a03969e2c upstream.

    ntb_transport_tx_free_entry() never returns 0 with the current
    calculation. If head == tail, then it would return qp->tx_max_entry.
    Change compare to tail >= head and when they are equal, a 0 would be
    returned.

    Fixes: e74bfeedad08 ("NTB: Add flow control to the ntb_netdev")
    Reviewed-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: renlonglong <ren.longlong@h3c.com>
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Jon Mason <jdmason@kudzu.us>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 446cb3d0a07c39be6bde93634843becd68753671
Author: Dave Jiang <dave.jiang@intel.com>
Date:   Tue Aug 22 09:04:45 2023 -0700

    ntb: Clean up tx tail index on link down
```

```
    commit cc79bd2738c2d40aba58b2be6ce47dc0e471df0e upstream.

    The tx tail index is not reset when the link goes down. This causes the
    tail index to go out of sync when the link goes down and comes back up.
    Refactor the ntb_qp_link_down_reset() and reset the tail index as well.

    Fixes: 2849b5d70641 ("NTB: Reset transport QP link stats on down")
    Reported-by: Yuan Y Lu <yuan.y.lu@intel.com>
    Tested-by: Yuan Y Lu <yuan.y.lu@intel.com>
    Reviewed-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Jon Mason <jdmason@kudzu.us>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 449b1978d4e50b14b2a4de3a26d0b074593926a3
Author: Dave Jiang <dave.jiang@intel.com>
Date:   Tue Aug 22 09:04:51 2023 -0700

    ntb: Drop packets when qp link is down

    commit f195a1a6fe416882984f8bd6c61afc1383171860 upstream.

    Currently when the transport receive packets after netdev has closed the
    transport returns error and triggers tx errors to be incremented and
    carrier to be stopped. There is no reason to return error if the device is
    already closed. Drop the packet and return 0.

    Fixes: e26a5843f7f5 ("NTB: Split ntb_hw_intel and ntb_transport drivers")
    Reported-by: Yuan Y Lu <yuan.y.lu@intel.com>
    Tested-by: Yuan Y Lu <yuan.y.lu@intel.com>
    Reviewed-by: Logan Gunthorpe <logang@deltatee.com>
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Jon Mason <jdmason@kudzu.us>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e79185c77c6e49720cabfd83611f5aed44369599
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Sun Aug 27 10:53:51 2023 +0200

    dt-bindings: PCI: qcom: Fix SDX65 compatible

    commit 15d63a897f79f465d71fb55cc11c6b7e20c19391 upstream.

    Commit c0aba9f32801 ("dt-bindings: PCI: qcom: Add SDX65 SoC") adding
    SDX65 was never tested and is clearly bogus.  The qcom,sdx65-pcie-ep
    compatible is followed by a fallback in DTS, and there is no driver
    matched by this compatible.  Driver matches by its fallback
    qcom,sdx55-pcie-ep.  This also fixes dtbs_check warnings like:

      qcom-sdx65-mtp.dtb: pcie-ep@1c00000: compatible: ['qcom,sdx65-pcie-ep', 'qcom,sdx55-
pcie-ep'] is too long

    [kwilczynski: commit log]
    Fixes: c0aba9f32801 ("dt-bindings: PCI: qcom: Add SDX65 SoC")
    Link: https://lore.kernel.org/linux-pci/20230827085351.21932-1-
krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Acked-by: Conor Dooley <conor.dooley@microchip.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3b89dbc8723e3941add8bb958ed483018cb4612a
Author: Feiyang Chen <chenfeiyang@loongson.cn>
Date:   Thu Aug 24 09:37:38 2023 +0800
```

```
    PCI/PM: Only read PCI_PM_CTRL register when available

    commit 5694ba13b004eea683c6d4faeb6d6e7a9636bda0 upstream.

    For a device with no Power Management Capability, pci_power_up() previously
    returned 0 (success) if the platform was able to put the device in D0,
    which led to pci_set_full_power_state() trying to read PCI_PM_CTRL, even
    though it doesn't exist.

    Since dev->pm_cap == 0 in this case, pci_set_full_power_state() actually
    read the wrong register, interpreted it as PCI_PM_CTRL, and corrupted
    dev->current_state.  This led to messages like this in some cases:

      pci 0000:01:00.0: Refused to change power state from D3hot to D0

    To prevent this, make pci_power_up() always return a negative failure code
    if the device lacks a Power Management Capability, even if non-PCI platform
    power management has been able to put the device in D0.  The failure will
    prevent pci_set_full_power_state() from trying to access PCI_PM_CTRL.

    Fixes: e200904b275c ("PCI/PM: Split pci_power_up()")
    Link: https://lore.kernel.org/r/20230824013738.1894965-1-chenfeiyang@loongson.cn
    Signed-off-by: Feiyang Chen <chenfeiyang@loongson.cn>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: "Rafael J. Wysocki" <rafael@kernel.org>
    Cc: stable@vger.kernel.org      # v5.19+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e32fc2168aa6b477290392ddbb73d95f012b050c
Author: Dexuan Cui <decui@microsoft.com>
Date:   Wed Aug 16 10:59:39 2023 -0700

    PCI: hv: Fix a crash in hv_pci_restore_msi_msg() during hibernation

    commit 04bbe863241a9be7d57fb4cf217ee4a72f480e70 upstream.

    When a Linux VM with an assigned PCI device runs on Hyper-V, if the PCI
    device driver is not loaded yet (i.e. MSI-X/MSI is not enabled on the
    device yet), doing a VM hibernation triggers a panic in
    hv_pci_restore_msi_msg() -> msi_lock_descs(&pdev->dev), because
    pdev->dev.msi.data is still NULL.

    Avoid the panic by checking if MSI-X/MSI is enabled.

    Link: https://lore.kernel.org/r/20230816175939.21566-1-decui@microsoft.com
    Fixes: dc2b453290c4 ("PCI: hv: Rework MSI handling")
    Signed-off-by: Dexuan Cui <decui@microsoft.com>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Reviewed-by: sathyanarayanan.kuppuswamy@linux.intel.com
    Reviewed-by: Michael Kelley <mikelley@microsoft.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a08713b9d9031683b83b3ecf12bad40a1ca35211
Author: Ross Lagerwall <ross.lagerwall@citrix.com>
Date:   Wed Sep 6 12:08:46 2023 +0100

    PCI: Free released resource after coalescing

    commit 8ec9c1d5d0a5a4744516adb483b97a238892f9d5 upstream.

    release_resource() doesn't actually free the resource or resource list
    entry so free the resource list entry to avoid a leak.
```

```
    Closes: https://lore.kernel.org/r/878r9sga1t.fsf@kernel.org/
    Fixes: e54223275ba1 ("PCI: Release resource invalidated by coalescing")
    Link: https://lore.kernel.org/r/20230906110846.225369-1-ross.lagerwall@citrix.com
    Reported-by: Kalle Valo <kvalo@kernel.org>
    Tested-by: Kalle Valo <kvalo@kernel.org>
    Signed-off-by: Ross Lagerwall <ross.lagerwall@citrix.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Cc: stable@vger.kernel.org      # v5.16+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8395ac8f296075aad4ccfe9826e4efd61d541583
Author: Ranjan Kumar <ranjan.kumar@broadcom.com>
Date:   Tue Aug 29 14:30:19 2023 +0530

    scsi: mpt3sas: Perform additional retries if doorbell read returns 0

    commit 4ca10f3e31745d35249a727ecd108eb58f0a8c5e upstream.

    The driver retries certain register reads 3 times if the returned value is
    0. This was done because the controller could return 0 for certain
    registers if other registers were being accessed concurrently by the BMC.

    In certain systems with increased BMC interactions, the register values
    returned can be 0 for longer than 3 retries. Change the retry count from 3
    to 30 for the affected registers to prevent problems with out-of-band
    management.

    Fixes: b899202901a8 ("scsi: mpt3sas: Add separate function for aero doorbell reads")
    Cc: stable@vger.kernel.org
    Signed-off-by: Ranjan Kumar <ranjan.kumar@broadcom.com>
    Link: https://lore.kernel.org/r/20230829090020.5417-2-ranjan.kumar@broadcom.com
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 5487df82bf0f71e286cc56ebb17e6141749d51d1
Author: Nilesh Javali <njavali@marvell.com>
Date:   Mon Aug 21 18:30:44 2023 +0530

    Revert "scsi: qla2xxx: Fix buffer overrun"

    commit 641671d97b9199f1ba35ccc2222d4b189a6a5de5 upstream.

    Revert due to Get PLOGI Template failed.
    This reverts commit b68710a8094fdffe8dd4f7a82c82649f479bb453.

    Cc: stable@vger.kernel.org
    Signed-off-by: Nilesh Javali <njavali@marvell.com>
    Link: https://lore.kernel.org/r/20230821130045.34850-9-njavali@marvell.com
    Reviewed-by: Himanshu Madhani <himanshu.madhani@oracle.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f0ea91d58b5a1101fb5c9c77442fa6537c51324d
Author: Yang Yingliang <yangyingliang@huawei.com>
Date:   Tue Jul 25 21:00:24 2023 +0800

    media: nxp: Fix wrong return pointer check in mxc_isi_crossbar_init()

    commit 4b60db99babad0254129ddc58e0927ffa9e93e35 upstream.

    It should check 'xbar->inputs', when allocate memory for it.

    Cc: stable@vger.kernel.org
    Fixes: cf21f328fcaf ("media: nxp: Add i.MX8 ISI driver")
    Signed-off-by: Yang Yingliang <yangyingliang@huawei.com>
```

```
    Reviewed-by: Fabio Estevam <festevam@gmail.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f67ef268eedeadb513583507f6c02203a58da8e1
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue May 30 14:30:36 2023 +0200

    media: venus: hfi_venus: Write to VIDC_CTRL_INIT after unmasking interrupts

    commit d74e481609808330b4625b3691cf01e1f56e255e upstream.

    The startup procedure shouldn't be started with interrupts masked, as that
    may entail silent failures.

    Kick off initialization only after the interrupts are unmasked.

    Cc: stable@vger.kernel.org # v4.12+
    Fixes: d96d3f30c0f2 ("[media] media: venus: hfi: add Venus HFI files")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Stanimir Varbanov <stanimir.k.varbanov@gmail.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a38f77781a0995100979e790dc66a774f2ae36ed
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Fri Sep 8 10:20:36 2023 +0100

    media: dvb: symbol fixup for dvb_attach()

    commit 86495af1171e1feec79faa9b64c05c89f46e41d1 upstream.

    In commit 9011e49d54dc ("modules: only allow symbol_get of
    EXPORT_SYMBOL_GPL modules") the use of symbol_get is properly restricted
    to GPL-only marked symbols.  This interacts oddly with the DVB logic
    which only uses dvb_attach() to load the dvb driver which then uses
    symbol_get().

    Fix this up by properly marking all of the dvb_attach attach symbols as
    EXPORT_SYMBOL_GPL().

    Fixes: 9011e49d54dc ("modules: only allow symbol_get of EXPORT_SYMBOL_GPL modules")
    Cc: stable <stable@kernel.org>
    Reported-by: Stefan Lippers-Hollmann <s.l-h@gmx.de>
    Cc: Mauro Carvalho Chehab <mchehab@kernel.org>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: linux-media@vger.kernel.org
    Cc: linux-modules@vger.kernel.org
    Acked-by: Luis Chamberlain <mcgrof@kernel.org>
    Acked-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Link: https://lore.kernel.org/r/20230908092035.3815268-2-gregkh@linuxfoundation.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6a73d9600896200639a342993c284194ec434a43
Author: Ding Xiang <dingxiang@cmss.chinamobile.com>
Date:   Wed Aug 30 18:11:48 2023 +0800

    selftests/landlock: Fix a resource leak

    commit 2a2015495142ee0a35711b5dcf7b215c31489f27 upstream.

    The opened file should be closed before return, otherwise resource leak
    will occur.
```

```
    Signed-off-by: Ding Xiang <dingxiang@cmss.chinamobile.com>
    Link: https://lore.kernel.org/r/20230830101148.3738-1-dingxiang@cmss.chinamobile.com
    Fixes: 3de64b656b3c ("selftests/landlock: Add supports_filesystem() helper")
    Signed-off-by: Mickaël Salaün <mic@digikod.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 976d377244541f753c8fbabccde8002b8bb2c428
Author: Vitaly Rodionov <vitalyr@opensource.cirrus.com>
Date:   Mon Sep 4 17:00:33 2023 +0100

    ALSA: hda/cirrus: Fix broken audio on hardware with two CS42L42 codecs.

    commit 99bf5b0baac941176a6a3d5cef7705b29808de34 upstream.

    Recently in v6.3-rc1 there was a change affecting behaviour of hrtimers
    (commit 0c52310f260014d95c1310364379772cb74cf82d) and causing
    few issues on platforms with two CS42L42 codecs. Canonical/Dell
    has reported an issue with Vostro-3910.
    We need to increase this value by 15ms.

    Link: https://bugs.launchpad.net/somerville/+bug/2031060
    Fixes: 9fb9fa18fb50 ("ALSA: hda/cirrus: Add extra 10 ms delay to allow PLL settle and
lock.")
    Signed-off-by: Vitaly Rodionov <vitalyr@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20230904160033.908135-1-vitalyr@opensource.cirrus.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 080efcc5685cfd41a064b730e1e3e8c25ddca28d
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Sep 5 10:12:10 2023 +0200

    ALSA: seq: Fix snd_seq_expand_var_event() call to user-space

    commit 86496fd4a2fabb7c978fdaca2d4b718207a96d36 upstream.

    The recent fix to clear the padding bytes at
    snd_seq_expand_var_event() broke the read to user-space with
    in_kernel=0 parameter.  For user-space address, it has to use
    clear_user() instead of memset().

    Fixes: f80e6d60d677 ("ALSA: seq: Clear padded bytes at expanding events")
    Reported-and-tested-by: Ash Holland <ash@sorrel.sh>
    Closes: https://lore.kernel.org/r/8a555319-9f31-4ea2-878f-adc338bc40d4@sorrel.sh
    Link: https://lore.kernel.org/r/20230905052631.18240-1-tiwai@suse.de
    Link: https://lore.kernel.org/r/20230905081210.6731-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f819b343aa95d24d5f7d6e06660c7f62591abc5f
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Sep 5 07:45:11 2023 +0200

    ALSA: usb-audio: Fix potential memory leaks at error path for UMP open

    commit b1757fa30ef14f254f4719bf6f7d54a4c8207216 upstream.

    The allocation and initialization errors at alloc_midi_urbs() that is
    called at MIDI 2.0 / UMP device are supposed to be handled at the
    caller side by invoking free_midi_urbs().  However, free_midi_urbs()
    loops only for ep->num_urbs entries, and since ep->num_entries wasn't
    updated yet at the allocation / init error in alloc_midi_urbs(), this
    entry won't be released.
```

```
    The intention of free_midi_urbs() is to release the whole elements, so
    change the loop size to NUM_URBS to scan over all elements for fixing
    the missed releases.

    Also, the call of free_midi_urbs() is missing at
    snd_usb_midi_v2_open().  Although it'll be released later at
    reopen/close or disconnection, it's better to release immediately at
    the error path.

    Fixes: ff49d1df79ae ("ALSA: usb-audio: USB MIDI 2.0 UMP support")
    Reported-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Closes:
https://lore.kernel.org/r/fc275ed315b9157952dcf2744ee7bdb78defdb5f.1693746347.git.christophe
.jaillet@wanadoo.fr
    Link: https://lore.kernel.org/r/20230905054511.20502-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit fcdf904e866de0e3715835e50409fda3b2590527
Author: Will Deacon <will@kernel.org>
Date:   Thu Sep 7 09:54:11 2023 +0100

    arm64: csum: Fix OoB access in IP checksum code for negative lengths

    commit 8bd795fedb8450ecbef18eeadbd23ed8fc7630f5 upstream.

    Although commit c2c24edb1d9c ("arm64: csum: Fix pathological zero-length
    calls") added an early return for zero-length input, syzkaller has
    popped up with an example of a _negative_ length which causes an
    undefined shift and an out-of-bounds read:

     | BUG: KASAN: slab-out-of-bounds in do_csum+0x44/0x254 arch/arm64/lib/csum.c:39
     | Read of size 4294966928 at addr ffff0000d7ac0170 by task syz-executor412/5975
     |
     | CPU: 0 PID: 5975 Comm: syz-executor412 Not tainted 6.4.0-rc4-syzkaller-g908f31f2a05b
#0
     | Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
05/25/2023
     | Call trace:
     |  dump_backtrace+0x1b8/0x1e4 arch/arm64/kernel/stacktrace.c:233
     |  show_stack+0x2c/0x44 arch/arm64/kernel/stacktrace.c:240
     |  __dump_stack lib/dump_stack.c:88 [inline]
     |  dump_stack_lvl+0xd0/0x124 lib/dump_stack.c:106
     |  print_address_description mm/kasan/report.c:351 [inline]
     |  print_report+0x174/0x514 mm/kasan/report.c:462
     |  kasan_report+0xd4/0x130 mm/kasan/report.c:572
     |  kasan_check_range+0x264/0x2a4 mm/kasan/generic.c:187
     |  __kasan_check_read+0x20/0x30 mm/kasan/shadow.c:31
     |  do_csum+0x44/0x254 arch/arm64/lib/csum.c:39
     |  csum_partial+0x30/0x58 lib/checksum.c:128
     |  gso_make_checksum include/linux/skbuff.h:4928 [inline]
     |  __udp_gso_segment+0xaf4/0x1bc4 net/ipv4/udp_offload.c:332
     |  udp6_ufo_fragment+0x540/0xca0 net/ipv6/udp_offload.c:47
     |  ipv6_gso_segment+0x5cc/0x1760 net/ipv6/ip6_offload.c:119
     |  skb_mac_gso_segment+0x2b4/0x5b0 net/core/gro.c:141
     |  __skb_gso_segment+0x250/0x3d0 net/core/dev.c:3401
     |  skb_gso_segment include/linux/netdevice.h:4859 [inline]
     |  validate_xmit_skb+0x364/0xdbc net/core/dev.c:3659
     |  validate_xmit_skb_list+0x94/0x130 net/core/dev.c:3709
     |  sch_direct_xmit+0xe8/0x548 net/sched/sch_generic.c:327
     |  __dev_xmit_skb net/core/dev.c:3805 [inline]
     |  __dev_queue_xmit+0x147c/0x3318 net/core/dev.c:4210
     |  dev_queue_xmit include/linux/netdevice.h:3085 [inline]
     |  packet_xmit+0x6c/0x318 net/packet/af_packet.c:276
     |  packet_snd net/packet/af_packet.c:3081 [inline]
```

```
      |  packet_sendmsg+0x376c/0x4c98 net/packet/af_packet.c:3113
      |  sock_sendmsg_nosec net/socket.c:724 [inline]
      |  sock_sendmsg net/socket.c:747 [inline]
      |  __sys_sendto+0x3b4/0x538 net/socket.c:2144

    Extend the early return to reject negative lengths as well, aligning our
    implementation with the generic code in lib/checksum.c

    Cc: Robin Murphy <robin.murphy@arm.com>
    Fixes: 5777eaed566a ("arm64: Implement optimised checksum routine")
    Reported-by: syzbot+4a9f9820bd8d302e22f7@syzkaller.appspotmail.com
    Link: https://lore.kernel.org/r/000000000000e0e94c0603f8d213@google.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d25033932cea34a0a097919a6ac43c03cdf1c1e6
Author: Gabriel Krisman Bertazi <krisman@suse.de>
Date:   Mon Aug 28 19:42:49 2023 -0400

    io_uring: Don't set affinity on a dying sqpoll thread

    commit bd6fc5da4c51107e1e0cec4a3a07963d1dae2c84 upstream.

    Syzbot reported a null-ptr-deref of sqd->thread inside
    io_sqpoll_wq_cpu_affinity.  It turns out the sqd->thread can go away
    from under us during io_uring_register, in case the process gets a
    fatal signal during io_uring_register.

    It is not particularly hard to hit the race, and while I am not sure
    this is the exact case hit by syzbot, it solves it.  Finally, checking
    ->thread is enough to close the race because we locked sqd while
    "parking" the thread, thus preventing it from going away.

    I reproduced it fairly consistently with a program that does:

    int main(void) {
      ...
      io_uring_queue_init(RING_LEN, &ring1, IORING_SETUP_SQPOLL);
      while (1) {
        io_uring_register_iowq_aff(ring, 1, &mask);
      }
    }

    Executed in a loop with timeout to trigger SIGTERM:
      while true; do timeout 1 /a.out ; done

    This will hit the following BUG() in very few attempts.

    BUG: kernel NULL pointer dereference, address: 00000000000007a8
    PGD 800000010e949067 P4D 800000010e949067 PUD 10e46e067 PMD 0
    Oops: 0000 [#1] PREEMPT SMP PTI
    CPU: 0 PID: 15715 Comm: dead-sqpoll Not tainted 6.5.0-rc7-next-20230825-g193296236fa0-
dirty #23
    Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 0.0.0 02/06/2015
    RIP: 0010:io_sqpoll_wq_cpu_affinity+0x27/0x70
    Code: 90 90 90 0f 1f 44 00 00 55 53 48 8b 9f 98 03 00 00 48 85 db 74 4f
    48 89 df 48 89 f5 e8 e2 f8 ff ff 48 8b 43 38 48 85 c0 74 22 <48> 8b b8
    a8 07 00 00 48 89 ee e8 ba b1 00 00 48 89 df 89 c5 e8 70
    RSP: 0018:ffffb04040ea7e70 EFLAGS: 00010282
    RAX: 0000000000000000 RBX: ffff93c010749e40 RCX: 0000000000000001
    RDX: 0000000000000000 RSI: ffffffffa7653331 RDI: 00000000ffffffff
    RBP: ffffb04040ea7eb8 R08: 0000000000000000 R09: c0000000ffffdfff
    R10: ffff93c01141b600 R11: ffffb04040ea7d18 R12: ffff93c00ea74840
    R13: 0000000000000011 R14: 0000000000000000 R15: ffff93c00ea74800
    FS:  00007fb7c276ab80(0000) GS:ffff93c36f200000(0000) knlGS:0000000000000000
```

```
      CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
      CR2: 00000000000007a8 CR3: 0000000111634003 CR4: 0000000000370ef0
      DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
      DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
      Call Trace:
       <TASK>
       ? __die_body+0x1a/0x60
       ? page_fault_oops+0x154/0x440
       ? do_user_addr_fault+0x174/0x7b0
       ? exc_page_fault+0x63/0x140
       ? asm_exc_page_fault+0x22/0x30
       ? io_sqpoll_wq_cpu_affinity+0x27/0x70
       __io_register_iowq_aff+0x2b/0x60
       __io_uring_register+0x614/0xa70
       __x64_sys_io_uring_register+0xaa/0x1a0
       do_syscall_64+0x3a/0x90
       entry_SYSCALL_64_after_hwframe+0x6e/0xd8
      RIP: 0033:0x7fb7c226fec9
      Code: 2e 00 b8 ca 00 00 00 0f 05 eb a5 66 0f 1f 44 00 00 48 89 f8 48 89
      f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01
      f0 ff ff 73 01 c3 48 8b 0d 97 7f 2d 00 f7 d8 64 89 01 48
      RSP: 002b:00007ffe2c0674f8 EFLAGS: 00000246 ORIG_RAX: 00000000000001ab
      RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 00007fb7c226fec9
      RDX: 00007ffe2c067530 RSI: 0000000000000011 RDI: 0000000000000003
      RBP: 00007ffe2c0675d0 R08: 00007ffe2c067550 R09: 00007ffe2c067550
      R10: 0000000000000001 R11: 0000000000000246 R12: 0000000000000000
      R13: 00007ffe2c067750 R14: 0000000000000000 R15: 0000000000000000
       </TASK>
      Modules linked in:
      CR2: 00000000000007a8
      ---[ end trace 0000000000000000 ]---

      Reported-by: syzbot+c74fea926a78b8a91042@syzkaller.appspotmail.com
      Fixes: ebdfefc09c6d ("io_uring/sqpoll: fix io-wq affinity when IORING_SETUP_SQPOLL is
  used")
      Signed-off-by: Gabriel Krisman Bertazi <krisman@suse.de>
      Link: https://lore.kernel.org/r/87v8cybuo6.fsf@suse.de
      Signed-off-by: Jens Axboe <axboe@kernel.dk>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 3ad272d8c9d46d44d5d6d2e5d37658caad7dd34b
  Author: Frank Li <Frank.Li@nxp.com>
  Date:   Thu Aug 31 10:13:24 2023 -0400

      i3c: master: svc: fix probe failure when no i3c device exist

      commit 6e13d6528be2f7e801af63c8153b87293f25d736 upstream.

      I3C masters are expected to support hot-join. This means at initialization
      time we might not yet discover any device and this should not be treated
      as a fatal error.

      During the DAA procedure which happens at probe time, if no device has
      joined, all CCC will be NACKed (from a bus perspective). This leads to an
      early return with an error code which fails the probe of the master.

      Let's avoid this by just telling the core through an I3C_ERROR_M2
      return command code that no device was discovered, which is a valid
      situation. This way the master will no longer bail out and fail to probe
      for a wrong reason.

      Cc: stable@vger.kernel.org
      Fixes: dd3c52846d59 ("i3c: master: svc: Add Silvaco I3C master driver")
      Signed-off-by: Frank Li <Frank.Li@nxp.com>
      Acked-by: Miquel Raynal <miquel.raynal@bootlin.com>
```

```
    Link: https://lore.kernel.org/r/20230831141324.2841525-1-Frank.Li@nxp.com
    Signed-off-by: Alexandre Belloni <alexandre.belloni@bootlin.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 357234fd65f17fc9c1ec1c3e68add2a665fb3ec7
Author: Naveen N Rao <naveen@kernel.org>
Date:   Mon Jun 19 15:17:19 2023 +0530

    powerpc/ftrace: Fix dropping weak symbols with older toolchains

    commit f6834c8c59a8e977a6f6e4f96c5d28dfa5db8430 upstream.

    The minimum level of gcc supported for building the kernel is v5.1.
    v5.x releases of gcc emitted a three instruction sequence for
    -mprofile-kernel:
            mflr    r0
            std     r0, 16(r1)
            bl      _mcount

    It is only with the v6.x releases that gcc started emitting the two
    instruction sequence for -mprofile-kernel, omitting the second store
    instruction.

    With the older three instruction sequence, the actual ftrace location
    can be the 5th instruction into a function. Update the allowed offset
    for ftrace location from 12 to 16 to accommodate the same.

    Cc: stable@vger.kernel.org
    Fixes: 7af82ff90a2b06 ("powerpc/ftrace: Ignore weak functions")
    Signed-off-by: Naveen N Rao <naveen@kernel.org>
    Reviewed-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
https://msgid.link/7b265908a9461e38fc756ef9b569703860a80621.1687166935.git.naveen@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1375c226a569afbf3c5b4b79c1fb4757b26964e3
Author: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
Date:   Wed Sep 6 12:08:16 2023 -0700

    powercap: intel_rapl: Fix invalid setting of Power Limit 4

    commit 081690e941188acfad41b8dbde2112029a2aa206 upstream.

    System runs at minimum performance, once powercap RAPL package domain
    enabled flag is changed from 1 to 0 to 1.

    Setting RAPL package domain enabled flag to 0, results in setting of
    power limit 4 (PL4) MSR 0x601 to 0. This implies disabling PL4 limit.
    The PL4 limit controls the peak power. So setting 0, results in some
    undesirable performance, which depends on hardware implementation.

    Even worse, when the enabled flag is set to 1 again. This will set PL4
    MSR value to 0x01, which means reduce peak power to 0.125W. This will
    force system to run at the lowest possible performance on every PL4
    supported system.

    Setting enabled flag should only affect the "enable" bit, not other
    bits. Here it is changing power limit.

    This is caused by a change which assumes that there is an enable bit in
    the PL4 MSR like other power limits. Although PL4 enable/disable bit is
    present with TPMI RAPL interface, it is not present with the MSR
    interface.
```

```
    There is a rapl_primitive_info defined for non existent PL4 enable bit
    and then it is used with the commit 9050a9cd5e4c ("powercap: intel_rapl:
    Cleanup Power Limits support") to enable PL4. This is wrong, hence remove
    this rapl primitive for PL4. Also in the function
    rapl_detect_powerlimit(), PL_ENABLE is used to check for the presence of
    power limits. Replace PL_ENABLE with PL_LIMIT, as PL_LIMIT must be
    present. Without this change, PL4 controls will not be available in the
    sysfs once rapl primitive for PL4 is removed.

    Fixes: 9050a9cd5e4c ("powercap: intel_rapl: Cleanup Power Limits support")
    Suggested-by: Zhang Rui <rui.zhang@intel.com>
    Signed-off-by: Srinivas Pandruvada <srinivas.pandruvada@linux.intel.com>
    Tested-by: Sumeet Pawnikar <sumeet.r.pawnikar@intel.com>
    Cc: 6.5+ <stable@vger.kernel.org> # 6.5+
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 77aaf22a9200b9557793c96debead911b80acc1c
Author: Hongchen Zhang <zhanghongchen@loongson.cn>
Date:   Wed Sep 6 22:53:09 2023 +0800

    LoongArch: mm: Add p?d_leaf() definitions

    commit 303be4b33562a5b689261ced1616bf16ad49efa7 upstream.

    When I do LTP test, LTP test case ksm06 caused panic at
            break_ksm_pmd_entry
              -> pmd_leaf (Huge page table but False)
              -> pte_present (panic)

    The reason is pmd_leaf() is not defined, So like commit 501b81046701
    ("mips: mm: add p?d_leaf() definitions") add p?d_leaf() definition for
    LoongArch.

    Fixes: 09cfefb7fa70 ("LoongArch: Add memory management")
    Cc: stable@vger.kernel.org
    Acked-by: David Hildenbrand <david@redhat.com>
    Signed-off-by: Hongchen Zhang <zhanghongchen@loongson.cn>
    Signed-off-by: Huacai Chen <chenhuacai@loongson.cn>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 971a486c19320fda47ca65b75e65458da0444426
Author: Max Filippov <jcmvbkbc@gmail.com>
Date:   Mon Jul 24 00:58:24 2023 -0700

    xtensa: PMU: fix base address for the newer hardware

    commit 687eb3c42f4ad81e7c947c50e2d865f692064291 upstream.

    With introduction of ERI access control in RG.0 base address of the PMU
    unit registers has changed. Add support for the new PMU configuration.

    Cc: stable@vger.kernel.org
    Signed-off-by: Max Filippov <jcmvbkbc@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f6740393ca37be2c32c7844cb0b5040c9de69657
Author: Hamza Mahfooz <hamza.mahfooz@amd.com>
Date:   Tue Aug 22 12:31:09 2023 -0400

    drm/amd/display: register edp_backlight_control() for DCN301

    commit 1611917f39bee1abfc01501238db8ac19649042d upstream.

    As made mention of in commit 099303e9a9bd ("drm/amd/display: eDP
```

```
    intermittent black screen during PnP"), we need to turn off the
    display's backlight before powering off an eDP display. Not doing so
    will result in undefined behaviour according to the eDP spec. So, set
    DCN301's edp_backlight_control() function pointer to
    dce110_edp_backlight_control().

    Cc: stable@vger.kernel.org
    Link: https://gitlab.freedesktop.org/drm/amd/-/issues/2765
    Fixes: 9c75891feef0 ("drm/amd/display: rework recent update PHY state commit")
    Suggested-by: Swapnil Patel <swapnil.patel@amd.com>
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6a94ac5b4c6db59aa61daceeeb6c6b55b31180b7
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Tue Jun 13 13:06:40 2023 +0200

    backlight/lv5207lp: Compare against struct fb_info.device

    commit 1ca8819320fd84e7d95b04e7668efc5f9fe9fa5c upstream.

    Struct lv5207lp_platform_data refers to a platform device within
    the Linux device hierarchy. The test in lv5207lp_backlight_check_fb()
    compares it against the fbdev device in struct fb_info.dev, which
    is different. Fix the test by comparing to struct fb_info.device.

    Fixes a bug in the backlight driver and prepares fbdev for making
    struct fb_info.dev optional.

    v2:
            * move renames into separate patch (Javier, Sam, Michael)

    Fixes: 82e5c40d88f9 ("backlight: Add Sanyo LV5207LP backlight driver")
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Cc: Yoshinori Sato <ysato@users.sourceforge.jp>
    Cc: Rich Felker <dalias@libc.org>
    Cc: John Paul Adrian Glaubitz <glaubitz@physik.fu-berlin.de>
    Cc: Lee Jones <lee@kernel.org>
    Cc: Daniel Thompson <daniel.thompson@linaro.org>
    Cc: Jingoo Han <jingoohan1@gmail.com>
    Cc: linux-sh@vger.kernel.org
    Cc: dri-devel@lists.freedesktop.org
    Cc: <stable@vger.kernel.org> # v3.12+
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230613110953.24176-6-
tzimmermann@suse.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 9772c42824b0dcbb50fb4d77f28ca8f3615a824a
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Tue Jun 13 13:06:36 2023 +0200

    backlight/bd6107: Compare against struct fb_info.device

    commit 992bdddaabfba19bdc77c1c7a4977b2aa41ec891 upstream.

    Struct bd6107_platform_data refers to a platform device within
    the Linux device hierarchy. The test in bd6107_backlight_check_fb()
    compares it against the fbdev device in struct fb_info.dev, which
    is different. Fix the test by comparing to struct fb_info.device.
```

```
    Fixes a bug in the backlight driver and prepares fbdev for making
    struct fb_info.dev optional.

    v2:
            * move renames into separate patch (Javier, Sam, Michael)

    Fixes: 67b43e590415 ("backlight: Add ROHM BD6107 backlight driver")
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Cc: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Cc: Lee Jones <lee@kernel.org>
    Cc: Daniel Thompson <daniel.thompson@linaro.org>
    Cc: Jingoo Han <jingoohan1@gmail.com>
    Cc: dri-devel@lists.freedesktop.org
    Cc: <stable@vger.kernel.org> # v3.12+
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230613110953.24176-2-
tzimmermann@suse.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f229b7aacb4ecc94a12ebf0910c2a3851a44795b
Author: Thomas Zimmermann <tzimmermann@suse.de>
Date:   Tue Jun 13 13:06:38 2023 +0200

    backlight/gpio_backlight: Compare against struct fb_info.device

    commit 7b91d017f77c1bda56f27c2f4bbb70de7c6eca08 upstream.

    Struct gpio_backlight_platform_data refers to a platform device within
    the Linux device hierarchy. The test in gpio_backlight_check_fb()
    compares it against the fbdev device in struct fb_info.dev, which
    is different. Fix the test by comparing to struct fb_info.device.

    Fixes a bug in the backlight driver and prepares fbdev for making
    struct fb_info.dev optional.

    v2:
            * move renames into separate patch (Javier, Sam, Michael)

    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Fixes: 8b770e3c9824 ("backlight: Add GPIO-based backlight driver")
    Cc: Laurent Pinchart <laurent.pinchart+renesas@ideasonboard.com>
    Cc: Rich Felker <dalias@libc.org>
    Cc: John Paul Adrian Glaubitz <glaubitz@physik.fu-berlin.de>
    Cc: Lee Jones <lee@kernel.org>
    Cc: Daniel Thompson <daniel.thompson@linaro.org>
    Cc: Jingoo Han <jingoohan1@gmail.com>
    Cc: linux-sh@vger.kernel.org
    Cc: dri-devel@lists.freedesktop.org
    Cc: <stable@vger.kernel.org> # v3.12+
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Reviewed-by: Daniel Thompson <daniel.thompson@linaro.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230613110953.24176-4-
tzimmermann@suse.de
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f16c01edebc5b174bd722da862ff140782ac0a1b
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Thu Sep 7 13:50:07 2023 +0100

    io_uring: break out of iowq iopoll on teardown

    commit 45500dc4e01c167ee063f3dcc22f51ced5b2b1e9 upstream.
```

```
    io-wq will retry iopoll even when it failed with -EAGAIN. If that
    races with task exit, which sets TIF_NOTIFY_SIGNAL for all its workers,
    such workers might potentially infinitely spin retrying iopoll again and
    again and each time failing on some allocation / waiting / etc. Don't
    keep spinning if io-wq is dying.

    Fixes: 561fb04a6a225 ("io_uring: replace workqueue usage with io-wq")
    Cc: stable@vger.kernel.org
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e88b673aeda6c7bac121c03ee293e9aad26fb693
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Fri Aug 11 13:53:41 2023 +0100

    io_uring/net: don't overflow multishot accept

    commit 1bfed23349716a7811645336a7ce42c4b8f250bc upstream.

    Don't allow overflowing multishot accept CQEs, we want to limit
    the grows of the overflow list.

    Cc: stable@vger.kernel.org
    Fixes: 4e86a2c980137 ("io_uring: implement multishot mode for accept")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/7d0d749649244873772623dd7747966f516fe6e2.1691757663.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1e2db9837be7d24a2a74eb3f3906d0872bee8907
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Fri Aug 11 13:53:42 2023 +0100

    io_uring/net: don't overflow multishot recv

    commit b2e74db55dd93d6db22a813c9a775b5dbf87c560 upstream.

    Don't allow overflowing multishot recv CQEs, it might get out of
    hand, hurt performance, and in the worst case scenario OOM the task.

    Cc: stable@vger.kernel.org
    Fixes: b3fdea6ecb55c ("io_uring: multishot recv")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/0b295634e8f1b71aa764c984608c22d85f88f75c.1691757663.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6aee0822a0a1972e7e363c28b2611658f16b931c
Author: Jens Axboe <axboe@kernel.dk>
Date:   Sun Aug 13 11:05:36 2023 -0600

    io_uring/sqpoll: fix io-wq affinity when IORING_SETUP_SQPOLL is used

    commit ebdfefc09c6de7897962769bd3e63a2ff443ebf5 upstream.

    If we setup the ring with SQPOLL, then that polling thread has its
    own io-wq setup. This means that if the application uses
    IORING_REGISTER_IOWQ_AFF to set the io-wq affinity, we should not be
    setting it for the invoking task, but rather the sqpoll task.
```

```
    Add an sqpoll helper that parks the thread and updates the affinity,
    and use that one if we're using SQPOLL.

    Fixes: fe76421d1da1 ("io_uring: allow user configurable IO thread CPU affinity")
    Cc: stable@vger.kernel.org # 5.10+
    Link: https://github.com/axboe/liburing/discussions/884
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c392179ed52880b72499778a2f141beb16d6fd9f
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Wed Aug 9 16:20:21 2023 +0100

    io_uring: break iopolling on signal

    commit dc314886cb3d0e4ab2858003e8de2917f8a3ccbd upstream.

    Don't keep spinning iopoll with a signal set. It'll eventually return
    back, e.g. by virtue of need_resched(), but it's not a nice user
    experience.

    Cc: stable@vger.kernel.org
    Fixes: def596e9557c9 ("io_uring: support for IO polling")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/eeba551e82cad12af30c3220125eb6cb244cc94c.1691594339.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0d2bb32800bae6ecfdae4bfae4b64bbb96c5b267
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Wed Aug 9 13:22:16 2023 +0100

    io_uring: fix false positive KASAN warnings

    commit 569f5308e54352a12181cc0185f848024c5443e8 upstream.

    io_req_local_work_add() peeks into the work list, which can be executed
    in the meanwhile. It's completely fine without KASAN as we're in an RCU
    read section and it's SLAB_TYPESAFE_BY_RCU. With KASAN though it may
    trigger a false positive warning because internal io_uring caches are
    sanitised.

    Remove sanitisation from the io_uring request cache for now.

    Cc: stable@vger.kernel.org
    Fixes: 8751d15426a31 ("io_uring: reduce scheduling due to tw")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/c6fbf7a82a341e66a0007c76eefd9d57f2d3ba51.1691541473.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 31cd0859dc1fe345ebacc319b72a03e9cf80cf41
Author: Matthew Wilcox (Oracle) <willy@infradead.org>
Date:   Wed Jul 26 22:58:17 2023 -0400

    XArray: Do not return sibling entries from xa_load()

    commit cbc02854331edc6dc22d8b77b6e22e38ebc7dd51 upstream.

    It is possible for xa_load() to observe a sibling entry pointing to
```

```
    another sibling entry.  An example:

    Thread A:               Thread B:
                            xa_store_range(xa, entry, 188, 191, gfp);
    xa_load(xa, 191);
    entry = xa_entry(xa, node, 63);
    [entry is a sibling of 188]
                            xa_store_range(xa, entry, 184, 191, gfp);
    if (xa_is_sibling(entry))
    offset = xa_to_sibling(entry);
    entry = xa_entry(xas->xa, node, offset);
    [entry is now a sibling of 184]

    It is sufficient to go around this loop until we hit a non-sibling entry.
    Sibling entries always point earlier in the node, so we are guaranteed
    to terminate this search.

    Signed-off-by: Matthew Wilcox (Oracle) <willy@infradead.org>
    Fixes: 6b24ca4a1a8d ("mm: Use multi-index entries in the page cache")
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1269b6d6813d118a463e69ce6fe739e256eeb196
Author: Alexandre Ghiti <alexghiti@rivosinc.com>
Date:   Tue Jul 4 09:43:56 2023 +0200

    riscv: Mark KASAN tmp* page tables variables as static

    commit dd7664d67b478afeb79a89e4586c2cd7707d17d6 upstream.

    tmp_pg_dir, tmp_p4d and tmp_pud are only used in kasan_init.c so they
    should be declared as static.

    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202306282202.bODptiGE-lkp@intel.com/
    Fixes: 96f9d4daf745 ("riscv: Rework kasan population functions")
    Signed-off-by: Alexandre Ghiti <alexghiti@rivosinc.com>
    Link: https://lore.kernel.org/r/20230704074357.233982-1-alexghiti@rivosinc.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4a9b900e3f43999925eaffa028af212a4cc9d096
Author: Alexandre Ghiti <alexghiti@rivosinc.com>
Date:   Tue Jul 4 09:43:57 2023 +0200

    riscv: Move create_tmp_mapping() to init sections

    commit 9bdd924803787ceeb10f1ea399e91d75fb05d3a7 upstream.

    This function is only used at boot time so mark it as __init.

    Fixes: 96f9d4daf745 ("riscv: Rework kasan population functions")
    Signed-off-by: Alexandre Ghiti <alexghiti@rivosinc.com>
    Link: https://lore.kernel.org/r/20230704074357.233982-2-alexghiti@rivosinc.com
    Cc: stable@vger.kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8bd9f307b9edd70bb85fe331ffde771fe70ad79c
Author: Gustavo A. R. Silva <gustavoars@kernel.org>
Date:   Wed Jun 7 22:12:11 2023 -0600

    ARM: OMAP2+: Fix -Warray-bounds warning in _pwrdm_state_switch()
```

```
    commit 847fb80cc01a54bc827b02547bb8743bdb59ddab upstream.

    If function pwrdm_read_prev_pwrst() returns -EINVAL, we will end
    up accessing array pwrdm->state_counter through negative index
    -22. This is wrong and the compiler is legitimately warning us
    about this potential problem.

    Fix this by sanity checking the value stored in variable _prev_
    before accessing array pwrdm->state_counter.

    Address the following -Warray-bounds warning:
    arch/arm/mach-omap2/powerdomain.c:178:45: warning: array subscript -22 is below array
bounds of 'unsigned int[4]' [-Warray-bounds]

    Link: https://github.com/KSPP/linux/issues/307
    Fixes: ba20bb126940 ("OMAP: PM counter infrastructure.")
    Cc: stable@vger.kernel.org
    Reported-by: kernel test robot <lkp@intel.com>
    Link: https://lore.kernel.org/lkml/20230607050639.LzbPn%25lkp@intel.com/
    Signed-off-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Message-ID: <ZIFVGwImU3kpaGeH@work>
    Acked-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7291af9a738d936c2d6869d030711dceb68404d0
Author: Yi Yang <yiyang13@huawei.com>
Date:   Thu Jun 29 20:33:28 2023 +0800

    ipmi_si: fix a memleak in try_smi_init()

    commit 6cf1a126de2992b4efe1c3c4d398f8de4aed6e3f upstream.

    Kmemleak reported the following leak info in try_smi_init():

    unreferenced object 0xffff00018ecf9400 (size 1024):
      comm "modprobe", pid 2707763, jiffies 4300851415 (age 773.308s)
      backtrace:
        [<000000004ca5b312>] __kmalloc+0x4b8/0x7b0
        [<00000000953b1072>] try_smi_init+0x148/0x5dc [ipmi_si]
        [<000000006460d325>] 0xffff800081b10148
        [<0000000039206ea5>] do_one_initcall+0x64/0x2a4
        [<00000000601399ce>] do_init_module+0x50/0x300
        [<000000003c12ba3c>] load_module+0x7a8/0x9e0
        [<00000000c246fffe>] __se_sys_init_module+0x104/0x180
        [<00000000eea99093>] __arm64_sys_init_module+0x24/0x30
        [<0000000021b1ef87>] el0_svc_common.constprop.0+0x94/0x250
        [<0000000070f4f8b7>] do_el0_svc+0x48/0xe0
        [<000000005a05337f>] el0_svc+0x24/0x3c
        [<000000005eb248d6>] el0_sync_handler+0x160/0x164
        [<0000000030a59039>] el0_sync+0x160/0x180

    The problem was that when an error occurred before handlers registration
    and after allocating `new_smi->si_sm`, the variable wouldn't be freed in
    the error handling afterwards since `shutdown_smi()` hadn't been
    registered yet. Fix it by adding a `kfree()` in the error handling path
    in `try_smi_init()`.

    Cc: stable@vger.kernel.org # 4.19+
    Fixes: 7960f18a5647 ("ipmi_si: Convert over to a shutdown handler")
    Signed-off-by: Yi Yang <yiyang13@huawei.com>
    Co-developed-by: GONG, Ruiqi <gongruiqi@huaweicloud.com>
    Signed-off-by: GONG, Ruiqi <gongruiqi@huaweicloud.com>
    Message-Id: <20230629123328.2402075-1-gongruiqi@huaweicloud.com>
    Signed-off-by: Corey Minyard <minyard@acm.org>
```

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

commit e196fe4d57e63b83454bec4a70c988b3b6b9a57e
Author: Rick Wertenbroek <rick.wertenbroek@gmail.com>
Date:   Mon Jul 3 10:58:45 2023 +0200

    PCI: rockchip: Use 64-bit mask on MSI 64-bit PCI address

    commit cdb50033dd6dfcf02ae3d4ee56bc1a9555be6d36 upstream.

    A 32-bit mask was used on the 64-bit PCI address used for mapping MSIs.
    This would result in the upper 32 bits being unintentionally zeroed and
    MSIs getting mapped to incorrect PCI addresses if the address had any
    of the upper bits set.

    Replace 32-bit mask by appropriate 64-bit mask.

    [kwilczynski: use GENMASK_ULL() over GENMASK() for 32-bit compatibility]
    Fixes: dc73ed0f1b8b ("PCI: rockchip: Fix window mapping and address translation for
endpoint")
    Closes: https://lore.kernel.org/linux-pci/8d19e5b7-8fa0-44a4-90e2-
9bb06f5eb694@moroto.mountain
    Link: https://lore.kernel.org/linux-pci/20230703085845.2052008-1-
rick.wertenbroek@gmail.com
    Reported-by: Dan Carpenter <dan.carpenter@linaro.org>
    Signed-off-by: Rick Wertenbroek <rick.wertenbroek@gmail.com>
    Signed-off-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Reviewed-by: Damien Le Moal <dlemoal@kernel.org>
    Cc: stable@vger.kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7e6ec01108777b3434003de322cbba6e7d05007a
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Thu Jun 15 10:29:07 2023 +0200

    media: i2c: Add a camera sensor top level menu

    commit 7d3c7d2a2914e10bec3b9cdacdadb8e1f65f715a upstream.

    Select V4L2_FWNODE and VIDEO_V4L2_SUBDEV_API for all sensor drivers. This
    also adds the options to drivers that don't specifically need them, these
    are still seldom used drivers using old APIs. The upside is that these
    should now all compile --- many drivers have had missing dependencies.

    The "menu" is replaced by selectable "menuconfig" to select the needed
    V4L2_FWNODE and VIDEO_V4L2_SUBDEV_API options.

    Also select MEDIA_CONTROLLER which VIDEO_V4L2_SUBDEV_API effectively
    depends on, and add the I2C dependency to the menu.

    Reported-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Cc: stable@vger.kernel.org # for >= 6.1
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d4ac0337576d02c9fc7d8464bcbd0d7317f0a539
Author: Sakari Ailus <sakari.ailus@linux.intel.com>
Date:   Sat Jul 29 20:59:25 2023 +0200

    media: i2c: ccs: Check rules is non-NULL

    commit 607bcc4213d998d051541d8f10b5bbb7d546c0be upstream.
```

```
    Fix the following smatch warning:

    drivers/media/i2c/ccs/ccs-data.c:524 ccs_data_parse_rules() warn: address
    of NULL pointer 'rules'

    The CCS static data rule parser does not check an if rule has been
    obtained before checking for other rule types (which depend on the if
    rule). In practice this means parsing invalid CCS static data could lead
    to dereferencing a NULL pointer.

    Reported-by: Hans Verkuil <hverkuil@xs4all.nl>
    Fixes: a6b396f410b1 ("media: ccs: Add CCS static data parser library")
    Cc: stable@vger.kernel.org # for 5.11 and up
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 07f9e0c9987bf1c4ef57611ad2f789ba68978102
Author: Thomas Gleixner <tglx@linutronix.de>
Date:   Wed Aug 23 10:47:02 2023 +0200

    cpu/hotplug: Prevent self deadlock on CPU hot-unplug

    commit 2b8272ff4a70b866106ae13c36be7ecbef5d5da2 upstream.

    Xiongfeng reported and debugged a self deadlock of the task which initiates
    and controls a CPU hot-unplug operation vs. the CFS bandwidth timer.

        CPU1                                                    CPU2

    T1 sets cfs_quota
       starts hrtimer cfs_bandwidth 'period_timer'
    T1 is migrated to CPU2
                                                            T1 initiates offlining of CPU1
    Hotplug operation starts
      ...
    'period_timer' expires and is re-enqueued on CPU1
      ...
    take_cpu_down()
      CPU1 shuts down and does not handle timers
      anymore. They have to be migrated in the
      post dead hotplug steps by the control task.

                                                            T1 runs the post dead offline operation
                                                            T1 is scheduled out
                                                            T1 waits for 'period_timer' to expire

    T1 waits there forever if it is scheduled out before it can execute the hrtimer
    offline callback hrtimers_dead_cpu().

    Cure this by delegating the hotplug control operation to a worker thread on
    an online CPU. This takes the initiating user space task, which might be
    affected by the bandwidth timer, completely out of the picture.

    Reported-by: Xiongfeng Wang <wangxiongfeng2@huawei.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Yu Liao <liaoyu15@huawei.com>
    Acked-by: Vincent Guittot <vincent.guittot@linaro.org>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/lkml/8e785777-03aa-99e1-d20e-e956f5685be6@huawei.com
    Link: https://lore.kernel.org/r/87h6oqdq0i.ffs@tglx
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ab8eda202ecf4d4f413a1d6e0e45ab6775b0ae2b
```

```
Author: Joel Fernandes (Google) <joel@joelfernandes.org>
Date:   Mon Sep 4 18:08:04 2023 +0000

    mm/vmalloc: add a safer version of find_vm_area() for debug

    commit 0818e739b5c061b0251c30152380600fb9b84c0c upstream.

    It is unsafe to dump vmalloc area information when trying to do so from
    some contexts.  Add a safer trylock version of the same function to do a
    best-effort VMA finding and use it from vmalloc_dump_obj().

    [applied test robot feedback on unused function fix.]
    [applied Uladzislau feedback on locking.]
    Link: https://lkml.kernel.org/r/20230904180806.1002832-1-joel@joelfernandes.org
    Fixes: 98f180837a89 ("mm: Make mem_dump_obj() handle vmalloc() memory")
    Signed-off-by: Joel Fernandes (Google) <joel@joelfernandes.org>
    Reviewed-by: Uladzislau Rezki (Sony) <urezki@gmail.com>
    Reported-by: Zhen Lei <thunder.leizhen@huaweicloud.com>
    Cc: Paul E. McKenney <paulmck@kernel.org>
    Cc: Zqiang <qiang.zhang1211@gmail.com>
    Cc: <stable@vger.kernel.org>
    Cc: Matthew Wilcox (Oracle) <willy@infradead.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 4bf234f3e06e7e195959d151284f84e3dc69d27b
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Fri Jul 21 09:01:32 2023 -0700

    scsi: core: Fix the scsi_set_resid() documentation

    commit f669b8a683e4ee26fa5cafe19d71cec1786b556a upstream.

    Because scsi_finish_command() subtracts the residual from the buffer
    length, residual overflows must not be reported. Reflect this in the SCSI
    documentation. See also commit 9237f04e12cc ("scsi: core: Fix
    scsi_get/set_resid() interface")

    Cc: Damien Le Moal <dlemoal@kernel.org>
    Cc: Hannes Reinecke <hare@suse.de>
    Cc: Douglas Gilbert <dgilbert@interlog.com>
    Cc: stable@vger.kernel.org
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Link: https://lore.kernel.org/r/20230721160154.874010-2-bvanassche@acm.org
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6ca39c0535573220739a98c648a5a99c5bd07d4a
Author: Kees Cook <keescook@chromium.org>
Date:   Thu Aug 10 22:45:32 2023 -0700

    printk: ringbuffer: Fix truncating buffer size min_t cast

    commit 53e9e33ede37a247d926db5e4a9e56b55204e66c upstream.

    If an output buffer size exceeded U16_MAX, the min_t(u16, ...) cast in
    copy_data() was causing writes to truncate. This manifested as output
    bytes being skipped, seen as %NUL bytes in pstore dumps when the available
    record size was larger than 65536. Fix the cast to no longer truncate
    the calculation.

    Cc: Petr Mladek <pmladek@suse.com>
    Cc: Sergey Senozhatsky <senozhatsky@chromium.org>
    Cc: Steven Rostedt <rostedt@goodmis.org>
    Cc: John Ogness <john.ogness@linutronix.de>
```

```
    Reported-by: Vijay Balakrishna <vijayb@linux.microsoft.com>
    Link: https://lore.kernel.org/lkml/d8bb1ec7-a4c5-43a2-9de0-
9643a70b899f@linux.microsoft.com/
    Fixes: b6cf8b3f3312 ("printk: add lockless ringbuffer")
    Cc: stable@vger.kernel.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Tested-by: Vijay Balakrishna <vijayb@linux.microsoft.com>
    Tested-by: Guilherme G. Piccoli <gpiccoli@igalia.com> # Steam Deck
    Reviewed-by: Tyler Hicks (Microsoft) <code@tyhicks.com>
    Tested-by: Tyler Hicks (Microsoft) <code@tyhicks.com>
    Reviewed-by: John Ogness <john.ogness@linutronix.de>
    Reviewed-by: Sergey Senozhatsky <senozhatsky@chromium.org>
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Signed-off-by: Petr Mladek <pmladek@suse.com>
    Link: https://lore.kernel.org/r/20230811054528.never.165-kees@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 8fb1601ec0a2c4c34fc2170af767e5c2a6400573
Author: Zqiang <qiang.zhang1211@gmail.com>
Date:   Mon Sep 4 18:08:05 2023 +0000

    rcu: dump vmalloc memory info safely

    commit c83ad36a18c02c0f51280b50272327807916987f upstream.

    Currently, for double invoke call_rcu(), will dump rcu_head objects memory
    info, if the objects is not allocated from the slab allocator, the
    vmalloc_dump_obj() will be invoke and the vmap_area_lock spinlock need to
    be held, since the call_rcu() can be invoked in interrupt context,
    therefore, there is a possibility of spinlock deadlock scenarios.

    And in Preempt-RT kernel, the rcutorture test also trigger the following
    lockdep warning:

    BUG: sleeping function called from invalid context at kernel/locking/spinlock_rt.c:48
    in_atomic(): 1, irqs_disabled(): 1, non_block: 0, pid: 1, name: swapper/0
    preempt_count: 1, expected: 0
    RCU nest depth: 1, expected: 1
    3 locks held by swapper/0/1:
     #0: ffffffffb534ee80 (fullstop_mutex){+.+.}-{4:4}, at: torture_init_begin+0x24/0xa0
     #1: ffffffffb5307940 (rcu_read_lock){....}-{1:3}, at: rcu_torture_init+0x1ec7/0x2370
     #2: ffffffffb536af40 (vmap_area_lock){+.+.}-{3:3}, at: find_vmap_area+0x1f/0x70
    irq event stamp: 565512
    hardirqs last  enabled at (565511): [<ffffffffb379b138>] __call_rcu_common+0x218/0x940
    hardirqs last disabled at (565512): [<ffffffffb5804262>] rcu_torture_init+0x20b2/0x2370
    softirqs last  enabled at (399112): [<ffffffffb36b2586>]
__local_bh_enable_ip+0x126/0x170
    softirqs last disabled at (399106): [<ffffffffb43fef59>] inet_register_protosw+0x9/0x1d0
    Preemption disabled at:
    [<ffffffffb58040c3>] rcu_torture_init+0x1f13/0x2370
    CPU: 0 PID: 1 Comm: swapper/0 Tainted: G        W          6.5.0-rc4-rt2-yocto-preempt-
rt+ #15
    Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.16.2-0-gea1b7a073390-
prebuilt.qemu.org 04/01/2014
    Call Trace:
     <TASK>
     dump_stack_lvl+0x68/0xb0
     dump_stack+0x14/0x20
     __might_resched+0x1aa/0x280
     ? __pfx_rcu_torture_err_cb+0x10/0x10
     rt_spin_lock+0x53/0x130
     ? find_vmap_area+0x1f/0x70
     find_vmap_area+0x1f/0x70
     vmalloc_dump_obj+0x20/0x60
     mem_dump_obj+0x22/0x90
```

```
       __call_rcu_common+0x5bf/0x940
       ? debug_smp_processor_id+0x1b/0x30
       call_rcu_hurry+0x14/0x20
       rcu_torture_init+0x1f82/0x2370
       ? __pfx_rcu_torture_leak_cb+0x10/0x10
       ? __pfx_rcu_torture_leak_cb+0x10/0x10
       ? __pfx_rcu_torture_init+0x10/0x10
       do_one_initcall+0x6c/0x300
       ? debug_smp_processor_id+0x1b/0x30
       kernel_init_freeable+0x2b9/0x540
       ? __pfx_kernel_init+0x10/0x10
       kernel_init+0x1f/0x150
       ret_from_fork+0x40/0x50
       ? __pfx_kernel_init+0x10/0x10
       ret_from_fork_asm+0x1b/0x30
       </TASK>

     The previous patch fixes this by using the deadlock-safe best-effort
     version of find_vm_area.  However, in case of failure print the fact that
     the pointer was a vmalloc pointer so that we print at least something.

     Link: https://lkml.kernel.org/r/20230904180806.1002832-2-joel@joelfernandes.org
     Fixes: 98f180837a89 ("mm: Make mem_dump_obj() handle vmalloc() memory")
     Signed-off-by: Zqiang <qiang.zhang1211@gmail.com>
     Signed-off-by: Joel Fernandes (Google) <joel@joelfernandes.org>
     Reported-by: Zhen Lei <thunder.leizhen@huaweicloud.com>
     Reviewed-by: Matthew Wilcox (Oracle) <willy@infradead.org>
     Cc: Paul E. McKenney <paulmck@kernel.org>
     Cc: Uladzislau Rezki (Sony) <urezki@gmail.com>
     Cc: <stable@vger.kernel.org>
     Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
     Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c7ab7e45ccef209809f8c2b00f497deec06b29c0
Author: Hou Tao <houtao1@huawei.com>
Date:   Thu Jul 13 21:54:13 2023 +0800

     virtio_pmem: add the missing REQ_OP_WRITE for flush bio

     commit c1dbd8a849183b9c12d257ad3043ecec50db50b3 upstream.

     When doing mkfs.xfs on a pmem device, the following warning was
     reported:

       ------------[ cut here ]------------
       WARNING: CPU: 2 PID: 384 at block/blk-core.c:751 submit_bio_noacct
       Modules linked in:
       CPU: 2 PID: 384 Comm: mkfs.xfs Not tainted 6.4.0-rc7+ #154
       Hardware name: QEMU Standard PC (i440FX + PIIX, 1996)
       RIP: 0010:submit_bio_noacct+0x340/0x520
       ......
       Call Trace:
        <TASK>
        ? submit_bio_noacct+0xd5/0x520
        submit_bio+0x37/0x60
        async_pmem_flush+0x79/0xa0
        nvdimm_flush+0x17/0x40
        pmem_submit_bio+0x370/0x390
        __submit_bio+0xbc/0x190
        submit_bio_noacct_nocheck+0x14d/0x370
        submit_bio_noacct+0x1ef/0x520
        submit_bio+0x55/0x60
        submit_bio_wait+0x5a/0xc0
        blkdev_issue_flush+0x44/0x60
```

```
    The root cause is that submit_bio_noacct() needs bio_op() is either
    WRITE or ZONE_APPEND for flush bio and async_pmem_flush() doesn't assign
    REQ_OP_WRITE when allocating flush bio, so submit_bio_noacct just fail
    the flush bio.

    Simply fix it by adding the missing REQ_OP_WRITE for flush bio. And we
    could fix the flush order issue and do flush optimization later.

    Cc: stable@vger.kernel.org # 6.3+
    Fixes: b4a6bb3a67aa ("block: add a sanity check for non-write flush/fua bios")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Chaitanya Kulkarni <kch@nvidia.com>
    Reviewed-by: Pankaj Gupta <pankaj.gupta@amd.com>
    Tested-by: Pankaj Gupta <pankaj.gupta@amd.com>
    Signed-off-by: Hou Tao <houtao1@huawei.com>
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d8c6fc85981821dd9908c5825c96b81d56296440
Author: Takashi Iwai <tiwai@suse.de>
Date:   Tue Aug 29 15:43:44 2023 +0200

    ALSA: pcm: Fix missing fixup call in compat hw_refine ioctl

    commit 358040e3807754944dbddf948a23c6d914297ed7 upstream.

    The update of rate_num/den and msbits were factored out to
    fixup_unreferenced_params() function to be called explicitly after the
    hw_refine or hw_params procedure.  It's called from
    snd_pcm_hw_refine_user(), but it's forgotten in the PCM compat ioctl.
    This ended up with the incomplete rate_num/den and msbits parameters
    when 32bit compat ioctl is used.

    This patch adds the missing call in snd_pcm_ioctl_hw_params_compat().

    Reported-by: Meng_Cai@novatek.com.cn
    Fixes: f9a076bff053 ("ALSA: pcm: calculate non-mask/non-interval parameters always when
possible")
    Reviewed-by: Takashi Sakamoto <o-takashi@sakamocchi.jp>
    Reviewed-by: Jaroslav Kysela <perex@perex.cz>
    Cc: <stable@vger.kernel.org>
    Link: https://lore.kernel.org/r/20230829134344.31588-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6b171ad315005a90541bce8c18748f1631cd4ff8
Author: Kalesh Singh <kaleshsingh@google.com>
Date:   Tue Aug 1 19:56:02 2023 -0700

    Multi-gen LRU: fix per-zone reclaim

    commit 669281ee7ef731fb5204df9d948669bf32a5e68d upstream.

    MGLRU has a LRU list for each zone for each type (anon/file) in each
    generation:

            long nr_pages[MAX_NR_GENS][ANON_AND_FILE][MAX_NR_ZONES];

    The min_seq (oldest generation) can progress independently for each
    type but the max_seq (youngest generation) is shared for both anon and
    file. This is to maintain a common frame of reference.

    In order for eviction to advance the min_seq of a type, all the per-zone
    lists in the oldest generation of that type must be empty.
```

```
    The eviction logic only considers pages from eligible zones for
    eviction or promotion.

        scan_folios() {
            ...
            for (zone = sc->reclaim_idx; zone >= 0; zone--)  {
                ...
                sort_folio();       // Promote
                ...
                isolate_folio();    // Evict
            }
            ...
        }

    Consider the system has the movable zone configured and default 4
    generations. The current state of the system is as shown below
    (only illustrating one type for simplicity):

    Type: ANON

            Zone    DMA32     Normal    Movable    Device

            Gen 0       0          0        4GB         0

            Gen 1       0        1GB        1MB         0

            Gen 2     1MB        4GB        1MB         0

            Gen 3     1MB        1MB        1MB         0

    Now consider there is a GFP_KERNEL allocation request (eligible zone
    index <= Normal), evict_folios() will return without doing any work
    since there are no pages to scan in the eligible zones of the oldest
    generation. Reclaim won't make progress until triggered from a ZONE_MOVABLE
    allocation request; which may not happen soon if there is a lot of free
    memory in the movable zone. This can lead to OOM kills, although there
    is 1GB pages in the Normal zone of Gen 1 that we have not yet tried to
    reclaim.

    This issue is not seen in the conventional active/inactive LRU since
    there are no per-zone lists.

    If there are no (not enough) folios to scan in the eligible zones, move
    folios from ineligible zone (zone_index > reclaim_index) to the next
    generation. This allows for the progression of min_seq and reclaiming
    from the next generation (Gen 1).

    Qualcomm, Mediatek and raspberrypi [1] discovered this issue independently.

    [1] https://github.com/raspberrypi/linux/issues/5395

    Link: https://lkml.kernel.org/r/20230802025606.346758-1-kaleshsingh@google.com
    Fixes: ac35a4902374 ("mm: multi-gen LRU: minimal implementation")
    Signed-off-by: Kalesh Singh <kaleshsingh@google.com>
    Reported-by: Charan Teja Kalla <quic_charante@quicinc.com>
    Reported-by: Lecopzer Chen <lecopzer.chen@mediatek.com>
    Tested-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
  [mediatek]
    Tested-by: Charan Teja Kalla <quic_charante@quicinc.com>
    Cc: Yu Zhao <yuzhao@google.com>
    Cc: Barry Song <baohua@kernel.org>
    Cc: Brian Geffon <bgeffon@google.com>
    Cc: Jan Alexander Steffens (heftig) <heftig@archlinux.org>
    Cc: Matthias Brugger <matthias.bgg@gmail.com>
    Cc: Oleksandr Natalenko <oleksandr@natalenko.name>
```

```
      Cc: Qi Zheng <zhengqi.arch@bytedance.com>
      Cc: Steven Barrett <steven@liquorix.net>
      Cc: Suleiman Souhlal <suleiman@google.com>
      Cc: Suren Baghdasaryan <surenb@google.com>
      Cc: Aneesh Kumar K V <aneesh.kumar@linux.ibm.com>
      Cc: <stable@vger.kernel.org>
      Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit 3354c401c68d70567d1ef25d12f4e22a7813a3c6
  Author: Boris Brezillon <boris.brezillon@collabora.com>
  Date:   Wed Aug 9 13:31:08 2023 +0200

      PM / devfreq: Fix leak in devfreq_dev_release()

      commit 5693d077595de721f9ddbf9d37f40e5409707dfe upstream.

      srcu_init_notifier_head() allocates resources that need to be released
      with a srcu_cleanup_notifier_head() call.

      Reported by kmemleak.

      Fixes: 0fe3a66410a3 ("PM / devfreq: Add new DEVFREQ_TRANSITION_NOTIFIER notifier")
      Cc: <stable@vger.kernel.org>
      Signed-off-by: Boris Brezillon <boris.brezillon@collabora.com>
      Reviewed-by: Dhruva Gole <d-gole@ti.com>
      Signed-off-by: Chanwoo Choi <cw00.choi@samsung.com>
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit b1b45a2a2055ba565d72f7f7e8a89b9865d98d35
  Author: Helge Deller <deller@gmx.de>
  Date:   Mon Aug 28 17:29:46 2023 +0200

      parisc: ccio-dma: Create private runway procfs root entry

      commit 77e0ddf097d6d4ceaf898e088b133b99e0a97fa0 upstream.

      Create an own procfs "runway" root entry for the CCIO driver.
      No need to share it with the sba_iommu driver, as only one
      of those busses can be active in one machine anyway.

      Signed-off-by: Helge Deller <deller@gmx.de>
      Reported-by: kernel test robot <lkp@intel.com>
      Fixes: 547259580dfa ("parisc: Move proc_mckinley_root and proc_runway_root to
  sba_iommu")
      Cc: <stable@vger.kernel.org> # v6.5
      Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

  commit dad92377fc6e012e481c6f4a75e32e48fe2777db
  Author: Radoslaw Tyl <radoslawx.tyl@intel.com>
  Date:   Thu Aug 24 13:46:19 2023 -0700

      igb: set max size RX buffer when store bad packet is enabled

      commit bb5ed01cd2428cd25b1c88a3a9cba87055eb289f upstream.

      Increase the RX buffer size to 3K when the SBP bit is on. The size of
      the RX buffer determines the number of pages allocated which may not
      be sufficient for receive frames larger than the set MTU size.

      Cc: stable@vger.kernel.org
      Fixes: 89eaefb61dc9 ("igb: Support RX-ALL feature flag.")
      Reported-by: Manfred Rudigier <manfred.rudigier@omicronenergy.com>
      Signed-off-by: Radoslaw Tyl <radoslawx.tyl@intel.com>
      Tested-by: Arpana Arland <arpanax.arland@intel.com> (A Contingent worker at Intel)
```

```
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit e763891c5ae08ecf2296fc6ad090d9321f7b8eb6
Author: Daniel Mack <daniel@zonque.org>
Date:   Fri Sep 1 14:24:24 2023 +0200

    gpio: zynq: restore zynq_gpio_irq_reqres/zynq_gpio_irq_relres callbacks

    commit 180b10bd160b014448366e5bc86e0558f8acb74f upstream.

    Commit f56914393537 ("gpio: zynq: fix zynqmp_gpio not an immutable chip
    warning") ditched the open-coded resource allocation handlers in favor
    of the generic ones. These generic handlers don't maintain the PM
    runtime anymore, which causes a regression in that level IRQs are no
    longer reported.

    Restore the original handlers to fix this.

    Signed-off-by: Daniel Mack <daniel@zonque.org>
    Fixes: f56914393537 ("gpio: zynq: fix zynqmp_gpio not an immutable chip warning")
    Cc: stable@kernel.org
    Signed-off-by: Bartosz Golaszewski <bartosz.golaszewski@linaro.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 7066517d00fada0f81845bd1f84af36f11a162d9
Author: Christoph Böhmwalder <christoph.boehmwalder@linbit.com>
Date:   Wed Sep 6 15:30:34 2023 +0200

    drbd: swap bvec_set_page len and offset

    commit 4b9c2edaf7282d60e069551b4b28abc2932cd3e3 upstream.

    bvec_set_page has the following signature:

    static inline void bvec_set_page(struct bio_vec *bv, struct page *page,
                    unsigned int len, unsigned int offset)

    However, the usage in DRBD swaps the len and offset parameters. This
    leads to a bvec with length=0 instead of the intended length=4096, which
    causes sock_sendmsg to return -EIO.

    This leaves DRBD unable to transmit any pages and thus completely
    broken.

    Swapping the parameters fixes the regression.

    Fixes: eeac7405c735 ("drbd: Use sendmsg(MSG_SPLICE_PAGES) rather than sendpage()")
    Reported-by: Serguei Ivantsov <manowar@gsc-game.com>
    Link:
https://lore.kernel.org/regressions/CAKH+VT3YLmAn0Y8=q37UTDShqxDLsqPcQ4hBMzY7HPn7zNx+RQ@mail
.gmail.com/
    Cc: stable@vger.kernel.org
    Signed-off-by: Christoph Böhmwalder <christoph.boehmwalder@linbit.com>
    Link: https://lore.kernel.org/r/20230906133034.948817-1-christoph.boehmwalder@linbit.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6c26ed3c6abe86ddab0510529000b970b05c9b40
Author: Mohamed Khalfella <mkhalfella@purestorage.com>
Date:   Thu Aug 31 02:17:02 2023 -0600

    skbuff: skb_segment, Call zero copy functions before using skbuff frags
```

```
    commit 2ea35288c83b3d501a88bc17f2df8f176b5cc96f upstream.

    Commit bf5c25d60861 ("skbuff: in skb_segment, call zerocopy functions
    once per nskb") added the call to zero copy functions in skb_segment().
    The change introduced a bug in skb_segment() because skb_orphan_frags()
    may possibly change the number of fragments or allocate new fragments
    altogether leaving nrfrags and frag to point to the old values. This can
    cause a panic with stacktrace like the one below.

    [  193.894380] BUG: kernel NULL pointer dereference, address: 00000000000000bc
    [  193.895273] CPU: 13 PID: 18164 Comm: vh-net-17428 Kdump: loaded Tainted: G
O       5.15.123+ #26
    [  193.903919] RIP: 0010:skb_segment+0xb0e/0x12f0
    [  194.021892] Call Trace:
    [  194.027422]  <TASK>
    [  194.072861]  tcp_gso_segment+0x107/0x540
    [  194.082031]  inet_gso_segment+0x15c/0x3d0
    [  194.090783]  skb_mac_gso_segment+0x9f/0x110
    [  194.095016]  __skb_gso_segment+0xc1/0x190
    [  194.103131]  netem_enqueue+0x290/0xb10 [sch_netem]
    [  194.107071]  dev_qdisc_enqueue+0x16/0x70
    [  194.110884]  __dev_queue_xmit+0x63b/0xb30
    [  194.121670]  bond_start_xmit+0x159/0x380 [bonding]
    [  194.128506]  dev_hard_start_xmit+0xc3/0x1e0
    [  194.131787]  __dev_queue_xmit+0x8a0/0xb30
    [  194.138225]  macvlan_start_xmit+0x4f/0x100 [macvlan]
    [  194.141477]  dev_hard_start_xmit+0xc3/0x1e0
    [  194.144622]  sch_direct_xmit+0xe3/0x280
    [  194.147748]  __dev_queue_xmit+0x54a/0xb30
    [  194.154131]  tap_get_user+0x2a8/0x9c0 [tap]
    [  194.157358]  tap_sendmsg+0x52/0x8e0 [tap]
    [  194.167049]  handle_tx_zerocopy+0x14e/0x4c0 [vhost_net]
    [  194.173631]  handle_tx+0xcd/0xe0 [vhost_net]
    [  194.176959]  vhost_worker+0x76/0xb0 [vhost]
    [  194.183667]  kthread+0x118/0x140
    [  194.190358]  ret_from_fork+0x1f/0x30
    [  194.193670]  </TASK>

    In this case calling skb_orphan_frags() updated nr_frags leaving nrfrags
    local variable in skb_segment() stale. This resulted in the code hitting
    i >= nrfrags prematurely and trying to move to next frag_skb using
    list_skb pointer, which was NULL, and caused kernel panic. Move the call
    to zero copy functions before using frags and nr_frags.

    Fixes: bf5c25d60861 ("skbuff: in skb_segment, call zerocopy functions once per nskb")
    Signed-off-by: Mohamed Khalfella <mkhalfella@purestorage.com>
    Reported-by: Amit Goyal <agoyal@purestorage.com>
    Cc: stable@vger.kernel.org
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 85ebbbe845823be6f8c04b4901da9a0a6f866283
Author: Wander Lairson Costa <wander@redhat.com>
Date:   Mon Aug 28 19:12:55 2023 -0300

    netfilter: xt_sctp: validate the flag_info count

    commit e99476497687ef9e850748fe6d232264f30bc8f9 upstream.

    sctp_mt_check doesn't validate the flag_count field. An attacker can
    take advantage of that to trigger a OOB read and leak memory
    information.

    Add the field validation in the checkentry function.
```

```
    Fixes: 2e4e6a17af35 ("[NETFILTER] x_tables: Abstraction layer for {ip,ip6,arp}_tables")
    Cc: stable@vger.kernel.org
    Reported-by: Lucas Leong <wmliang@infosec.exchange>
    Signed-off-by: Wander Lairson Costa <wander@redhat.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 799cc0fb184408f688b030ea381844b16d1d9c62
Author: Wander Lairson Costa <wander@redhat.com>
Date:   Mon Aug 28 10:21:07 2023 -0300

    netfilter: xt_u32: validate user space input

    commit 69c5d284f67089b4750d28ff6ac6f52ec224b330 upstream.

    The xt_u32 module doesn't validate the fields in the xt_u32 structure.
    An attacker may take advantage of this to trigger an OOB read by setting
    the size fields with a value beyond the arrays boundaries.

    Add a checkentry function to validate the structure.

    This was originally reported by the ZDI project (ZDI-CAN-18408).

    Fixes: 1b50b8a371e9 ("[NETFILTER]: Add u32 match")
    Cc: stable@vger.kernel.org
    Signed-off-by: Wander Lairson Costa <wander@redhat.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit c5a0c560aeb4c36d7ae87e65f1eb94f2efebcbd7
Author: Xiao Liang <shaw.leon@gmail.com>
Date:   Fri Aug 25 13:33:27 2023 +0800

    netfilter: nft_exthdr: Fix non-linear header modification

    commit 28427f368f0e08d504ed06e74bc7cc79d6d06511 upstream.

    Fix skb_ensure_writable() size. Don't use nft_tcp_header_pointer() to
    make it explicit that pointers point to the packet (not local buffer).

    Fixes: 99d1712bc41c ("netfilter: exthdr: tcp option set support")
    Fixes: 7890cbea66e7 ("netfilter: exthdr: add support for tcp option removal")
    Cc: stable@vger.kernel.org
    Signed-off-by: Xiao Liang <shaw.leon@gmail.com>
    Signed-off-by: Pablo Neira Ayuso <pablo@netfilter.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit d95c8420efe684b964e3aa28108e9a354bcd7225
Author: Kyle Zeng <zengyhkyle@gmail.com>
Date:   Tue Sep 5 15:04:09 2023 -0700

    netfilter: ipset: add the missing IP_SET_HASH_WITH_NET0 macro for
ip_set_hash_netportnet.c

    commit 050d91c03b28ca479df13dfb02bcd2c60dd6a878 upstream.

    The missing IP_SET_HASH_WITH_NET0 macro in ip_set_hash_netportnet can
    lead to the use of wrong `CIDR_POS(c)` for calculating array offsets,
    which can lead to integer underflow. As a result, it leads to slab
    out-of-bound access.
    This patch adds back the IP_SET_HASH_WITH_NET0 macro to
    ip_set_hash_netportnet to address the issue.

    Fixes: 886503f34d63 ("netfilter: ipset: actually allow allowable CIDR 0 in
```

```
hash:net,port,net")
    Suggested-by: Jozsef Kadlecsik <kadlec@netfilter.org>
    Signed-off-by: Kyle Zeng <zengyhkyle@gmail.com>
    Acked-by: Jozsef Kadlecsik <kadlec@netfilter.org>
    Signed-off-by: Florian Westphal <fw@strlen.de>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 3138192865c2a1f089dd27a7d80a7271ecd468e7
Author: Eric Dumazet <edumazet@google.com>
Date:   Tue Sep 5 04:23:38 2023 +0000

    igmp: limit igmpv3_newpack() packet size to IP_MAX_MTU

    commit c3b704d4a4a265660e665df51b129e8425216ed1 upstream.

    This is a follow up of commit 915d975b2ffa ("net: deal with integer
    overflows in kmalloc_reserve()") based on David Laight feedback.

    Back in 2010, I failed to realize malicious users could set dev->mtu
    to arbitrary values. This mtu has been since limited to 0x7fffffff but
    regardless of how big dev->mtu is, it makes no sense for igmpv3_newpack()
    to allocate more than IP_MAX_MTU and risk various skb fields overflows.

    Fixes: 57e1ab6eaddc ("igmp: refine skb allocations")
    Link: https://lore.kernel.org/netdev/d273628df80f45428e739274ab9ecb72@AcuMS.aculab.com/
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Reported-by: David Laight <David.Laight@ACULAB.COM>
    Cc: Kyle Zeng <zengyhkyle@gmail.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bf7da02d2b8faf324206e1cbe64a4813ff903cc1
Author: Eric Dumazet <edumazet@google.com>
Date:   Thu Aug 31 18:37:50 2023 +0000

    net: deal with integer overflows in kmalloc_reserve()

    commit 915d975b2ffa58a14bfcf16fafe00c41315949ff upstream.

    Blamed commit changed:
        ptr = kmalloc(size);
        if (ptr)
          size = ksize(ptr);

    to:
        size = kmalloc_size_roundup(size);
        ptr = kmalloc(size);

    This allowed various crash as reported by syzbot [1]
    and Kyle Zeng.

    Problem is that if @size is bigger than 0x80000001,
    kmalloc_size_roundup(size) returns 2^32.

    kmalloc_reserve() uses a 32bit variable (obj_size),
    so 2^32 is truncated to 0.

    kmalloc(0) returns ZERO_SIZE_PTR which is not handled by
    skb allocations.

    Following trace can be triggered if a netdev->mtu is set
    close to 0x7fffffff

    We might in the future limit netdev->mtu to more sensible
```

```
    limit (like KMALLOC_MAX_SIZE).

    This patch is based on a syzbot report, and also a report
    and tentative fix from Kyle Zeng.

    [1]
    BUG: KASAN: user-memory-access in __build_skb_around net/core/skbuff.c:294 [inline]
    BUG: KASAN: user-memory-access in __alloc_skb+0x3c4/0x6e8 net/core/skbuff.c:527
    Write of size 32 at addr 00000000fffffd10 by task syz-executor.4/22554

    CPU: 1 PID: 22554 Comm: syz-executor.4 Not tainted 6.1.39-syzkaller #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
07/03/2023
    Call trace:
    dump_backtrace+0x1c8/0x1f4 arch/arm64/kernel/stacktrace.c:279
    show_stack+0x2c/0x3c arch/arm64/kernel/stacktrace.c:286
    __dump_stack lib/dump_stack.c:88 [inline]
    dump_stack_lvl+0x120/0x1a0 lib/dump_stack.c:106
    print_report+0xe4/0x4b4 mm/kasan/report.c:398
    kasan_report+0x150/0x1ac mm/kasan/report.c:495
    kasan_check_range+0x264/0x2a4 mm/kasan/generic.c:189
    memset+0x40/0x70 mm/kasan/shadow.c:44
    __build_skb_around net/core/skbuff.c:294 [inline]
    __alloc_skb+0x3c4/0x6e8 net/core/skbuff.c:527
    alloc_skb include/linux/skbuff.h:1316 [inline]
    igmpv3_newpack+0x104/0x1088 net/ipv4/igmp.c:359
    add_grec+0x81c/0x1124 net/ipv4/igmp.c:534
    igmpv3_send_cr net/ipv4/igmp.c:667 [inline]
    igmp_ifc_timer_expire+0x1b0/0x1008 net/ipv4/igmp.c:810
    call_timer_fn+0x1c0/0x9f0 kernel/time/timer.c:1474
    expire_timers kernel/time/timer.c:1519 [inline]
    __run_timers+0x54c/0x710 kernel/time/timer.c:1790
    run_timer_softirq+0x28/0x4c kernel/time/timer.c:1803
    _stext+0x380/0xfbc
    ____do_softirq+0x14/0x20 arch/arm64/kernel/irq.c:79
    call_on_irq_stack+0x24/0x4c arch/arm64/kernel/entry.S:891
    do_softirq_own_stack+0x20/0x2c arch/arm64/kernel/irq.c:84
    invoke_softirq kernel/softirq.c:437 [inline]
    __irq_exit_rcu+0x1c0/0x4cc kernel/softirq.c:683
    irq_exit_rcu+0x14/0x78 kernel/softirq.c:695
    el0_interrupt+0x7c/0x2e0 arch/arm64/kernel/entry-common.c:717
    __el0_irq_handler_common+0x18/0x24 arch/arm64/kernel/entry-common.c:724
    el0t_64_irq_handler+0x10/0x1c arch/arm64/kernel/entry-common.c:729
    el0t_64_irq+0x1a0/0x1a4 arch/arm64/kernel/entry.S:584

    Fixes: 12d6c1d3a2ad ("skbuff: Proactively round up to kmalloc bucket size")
    Reported-by: syzbot <syzkaller@googlegroups.com>
    Reported-by: Kyle Zeng <zengyhkyle@gmail.com>
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Vlastimil Babka <vbabka@suse.cz>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 6eedcecf3a508e1d162bc6d000eea499244adf76
Author: Yuan Yao <yuanyaogoog@chromium.org>
Date:   Tue Aug 8 05:10:59 2023 +0000

    virtio_ring: fix avail_wrap_counter in virtqueue_add_packed

    [ Upstream commit 1acfe2c1225899eab5ab724c91b7e1eb2881b9ab ]

    In current packed virtqueue implementation, the avail_wrap_counter won't
    flip, in the case when the driver supplies a descriptor chain with a
    length equals to the queue size; total_sg == vq->packed.vring.num.
```

```
    Let’s assume the following situation:
    vq->packed.vring.num=4
    vq->packed.next_avail_idx: 1
    vq->packed.avail_wrap_counter: 0

    Then the driver adds a descriptor chain containing 4 descriptors.

    We expect the following result with avail_wrap_counter flipped:
    vq->packed.next_avail_idx: 1
    vq->packed.avail_wrap_counter: 1

    But, the current implementation gives the following result:
    vq->packed.next_avail_idx: 1
    vq->packed.avail_wrap_counter: 0

    To reproduce the bug, you can set a packed queue size as small as
    possible, so that the driver is more likely to provide a descriptor
    chain with a length equal to the packed queue size. For example, in
    qemu run following commands:
    sudo qemu-system-x86_64 \
    -enable-kvm \
    -nographic \
    -kernel "path/to/kernel_image" \
    -m 1G \
    -drive file="path/to/rootfs",if=none,id=disk \
    -device virtio-blk,drive=disk \
    -drive file="path/to/disk_image",if=none,id=rwdisk \
    -device virtio-blk,drive=rwdisk,packed=on,queue-size=4,\
    indirect_desc=off \
    -append "console=ttyS0 root=/dev/vda rw init=/bin/bash"

    Inside the VM, create a directory and mount the rwdisk device on it. The
    rwdisk will hang and mount operation will not complete.

    This commit fixes the wrap counter error by flipping the
    packed.avail_wrap_counter, when start of descriptor chain equals to the
    end of descriptor chain (head == i).

    Fixes: 1ce9e6055fa0 ("virtio_ring: introduce packed ring support")
    Signed-off-by: Yuan Yao <yuanyaogoog@chromium.org>
    Message-Id: <20230808051110.3492693-1-yuanyaogoog@chromium.org>
    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 628b53fc66ca1910a3cb53c3c7e44e59750c3668
Author: Jason Wang <jasowang@redhat.com>
Date:   Fri Aug 11 05:15:39 2023 -0400

    virtio_vdpa: build affinity masks conditionally

    [ Upstream commit ae15aceaa98ad9499763923f7890e345d9f46b60 ]

    We try to build affinity mask via create_affinity_masks()
    unconditionally which may lead several issues:

    - the affinity mask is not used for parent without affinity support
      (only VDUSE support the affinity now)
    - the logic of create_affinity_masks() might not work for devices
      other than block. For example it's not rare in the networking device
      where the number of queues could exceed the number of CPUs. Such
      case breaks the current affinity logic which is based on
      group_cpus_evenly() who assumes the number of CPUs are not less than
      the number of groups. This can trigger a warning[1]:
```

```
             if (ret >= 0)
                     WARN_ON(nr_present + nr_others < numgrps);
```

    Fixing this by only build the affinity masks only when

    - Driver passes affinity descriptor, driver like virtio-blk can make
      sure to limit the number of queues when it exceeds the number of CPUs
    - Parent support affinity setting config ops

    This help to avoid the warning. More optimizations could be done on
    top.

    [1]
    [682.146655] WARNING: CPU: 6 PID: 1550 at lib/group_cpus.c:400
group_cpus_evenly+0x1aa/0x1c0
    [682.146668] CPU: 6 PID: 1550 Comm: vdpa Not tainted 6.5.0-rc5jason+ #79
    [682.146671] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.2-0-
gea1b7a073390-prebuilt.qemu.org 04/01/2014
    [682.146673] RIP: 0010:group_cpus_evenly+0x1aa/0x1c0
    [682.146676] Code: 4c 89 e0 5b 5d 41 5c 41 5d 41 5e c3 cc cc cc cc e8 1b c4 74 ff 48
89 ef e8 13 ac 98 ff 4c 89 e7 45 31 e4 e8 08 ac 98 ff eb c2 <0f> 0b eb b6 e8 fd 05 c3 00 45
31 e4 eb e5 cc cc cc cc cc cc cc cc
    [682.146679] RSP: 0018:ffffc9000215f498 EFLAGS: 00010293
    [682.146682] RAX: 000000000001f1e0 RBX: 0000000000000041 RCX: 0000000000000000
    [682.146684] RDX: ffff888109922058 RSI: 0000000000000041 RDI: 0000000000000030
    [682.146686] RBP: ffff888109922058 R08: ffffc9000215f498 R09: ffffc9000215f4a0
    [682.146687] R10: 00000000000198d0 R11: 0000000000000030 R12: ffff888107e02800
    [682.146689] R13: 0000000000000030 R14: 0000000000000030 R15: 0000000000000041
    [682.146692] FS: 00007fef52315740(0000) GS:ffff888237380000(0000)
knlGS:0000000000000000
    [682.146695] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    [682.146696] CR2: 00007fef52509000 CR3: 0000000110dbc004 CR4: 0000000000370ee0
    [682.146698] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    [682.146700] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    [682.146701] Call Trace:
    [682.146703] <TASK>
    [682.146705] ? __warn+0x7b/0x130
    [682.146709] ? group_cpus_evenly+0x1aa/0x1c0
    [682.146712] ? report_bug+0x1c8/0x1e0
    [682.146717] ? handle_bug+0x3c/0x70
    [682.146721] ? exc_invalid_op+0x14/0x70
    [682.146723] ? asm_exc_invalid_op+0x16/0x20
    [682.146727] ? group_cpus_evenly+0x1aa/0x1c0
    [682.146729] ? group_cpus_evenly+0x15c/0x1c0
    [682.146731] create_affinity_masks+0xaf/0x1a0
    [682.146735] virtio_vdpa_find_vqs+0x83/0x1d0
    [682.146738] ? __pfx_default_calc_sets+0x10/0x10
    [682.146742] virtnet_find_vqs+0x1f0/0x370
    [682.146747] virtnet_probe+0x501/0xcd0
    [682.146749] ? vp_modern_get_status+0x12/0x20
    [682.146751] ? get_cap_addr.isra.0+0x10/0xc0
    [682.146754] virtio_dev_probe+0x1af/0x260
    [682.146759] really_probe+0x1a5/0x410

    Fixes: 3dad56823b53 ("virtio-vdpa: Support interrupt affinity spreading mechanism")
    Signed-off-by: Jason Wang <jasowang@redhat.com>
    Message-Id: <20230811091539.1359865-1-jasowang@redhat.com>
    Signed-off-by: Michael S. Tsirkin <mst@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 55f8a41525d57be4acf7b507f8ae562f0d92db8a
Author: Liao Chang <liaochang1@huawei.com>
Date: Tue Aug 29 07:03:18 2023 +0000

```
    cpufreq: Fix the race condition while updating the transition_task of policy

    [ Upstream commit 61bfbf7951ba561dcbdd5357702d3cbc2d447812 ]

    The field 'transition_task' of policy structure is used to track the
    task which is performing the frequency transition. Using this field to
    print a warning once detect a case where the same task is calling
    _begin() again before completing the preivous frequency transition via
    the _end().

    However, there is a potential race condition in _end() and _begin() APIs
    while updating the field 'transition_task' of policy, the scenario is
    depicted below:

                 Task A                              Task B

            /* 1st freq transition */
            Invoke _begin() {
                    ...
                    ...
            }
                                            /* 2nd freq transition */
                                            Invoke _begin() {
                                                    ... //waiting for A to
                                                    ... //clear
                                                    ... //transition_ongoing
                                                    ... //in _end() for
                                                    ... //the 1st transition
                                                            |
            Change the frequency                            |
                                                            |
            Invoke _end() {                                 |
                    ...                                     |
                    ...                                     |
                    transition_ongoing = false;             V
                                                    transition_ongoing = true;
                                                    transition_task = current;
                    transition_task = NULL;
                    ... //A overwrites the task
                    ... //performing the transition
                    ... //result in error warning.
            }

    To fix this race condition, the transition_lock of policy structure is
    now acquired before updating policy structure in _end() API. Which ensure
    that only one task can update the 'transition_task' field at a time.

    Link: https://lore.kernel.org/all/b3c61d8a-d52d-3136-fbf0-d1de9f1ba411@huawei.com/
    Fixes: ca654dc3a93d ("cpufreq: Catch double invocations of
cpufreq_freq_transition_begin/end")
    Signed-off-by: Liao Chang <liaochang1@huawei.com>
    Acked-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e942b92d8b6973955ae35b3d2e659bf4a3d557ba
Author: Vincent Whitchurch <vincent.whitchurch@axis.com>
Date:   Wed Aug 23 12:40:44 2023 +0200

    um: virt-pci: fix missing declaration warning

    [ Upstream commit 974b808d85abbc03c3914af63d60d5816aabf2ca ]

    Fix this warning which appears with W=1 and without CONFIG_OF:
```

```
      warning: no previous declaration for 'pcibios_get_phb_of_node'

    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202308230949.PphIIlhq-lkp@intel.com/
    Fixes: 314a1408b79a ("um: virt-pci: implement pcibios_get_phb_of_node()")
    Signed-off-by: Vincent Whitchurch <vincent.whitchurch@axis.com>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64f09d45e94547fbf219f36d1d02ac42742c028c
Author: Maciej S. Szmigiero <maciej.szmigiero@oracle.com>
Date:   Wed Aug 9 20:40:18 2023 +0200

    Drivers: hv: vmbus: Don't dereference ACPI root object handle

    [ Upstream commit 78e04bbff849b51b56f5925b1945db2c6e128b61 ]

    Since the commit referenced in the Fixes: tag below the VMBus client driver
    is walking the ACPI namespace up from the VMBus ACPI device to the ACPI
    namespace root object trying to find Hyper-V MMIO ranges.

    However, if it is not able to find them it ends trying to walk resources of
    the ACPI namespace root object itself.
    This object has all-ones handle, which causes a NULL pointer dereference
    in the ACPI code (from dereferencing this pointer with an offset).

    This in turn causes an oops on boot with VMBus host implementations that do
    not provide Hyper-V MMIO ranges in their VMBus ACPI device or its
    ancestors.
    The QEMU VMBus implementation is an example of such implementation.

    I guess providing these ranges is optional, since all tested Windows
    versions seem to be able to use VMBus devices without them.

    Fix this by explicitly terminating the lookup at the ACPI namespace root
    object.

    Note that Linux guests under KVM/QEMU do not use the Hyper-V PV interface
    by default - they only do so if the KVM PV interface is missing or
    disabled.

    Example stack trace of such oops:
    [ 3.710827] ? __die+0x1f/0x60
    [ 3.715030] ? page_fault_oops+0x159/0x460
    [ 3.716008] ? exc_page_fault+0x73/0x170
    [ 3.716959] ? asm_exc_page_fault+0x22/0x30
    [ 3.717957] ? acpi_ns_lookup+0x7a/0x4b0
    [ 3.718898] ? acpi_ns_internalize_name+0x79/0xc0
    [ 3.720018] acpi_ns_get_node_unlocked+0xb5/0xe0
    [ 3.721120] ? acpi_ns_check_object_type+0xfe/0x200
    [ 3.722285] ? acpi_rs_convert_aml_to_resource+0x37/0x6e0
    [ 3.723559] ? down_timeout+0x3a/0x60
    [ 3.724455] ? acpi_ns_get_node+0x3a/0x60
    [ 3.725412] acpi_ns_get_node+0x3a/0x60
    [ 3.726335] acpi_ns_evaluate+0x1c3/0x2c0
    [ 3.727295] acpi_ut_evaluate_object+0x64/0x1b0
    [ 3.728400] acpi_rs_get_method_data+0x2b/0x70
    [ 3.729476] ? vmbus_platform_driver_probe+0x1d0/0x1d0 [hv_vmbus]
    [ 3.730940] ? vmbus_platform_driver_probe+0x1d0/0x1d0 [hv_vmbus]
    [ 3.732411] acpi_walk_resources+0x78/0xd0
    [ 3.733398] vmbus_platform_driver_probe+0x9f/0x1d0 [hv_vmbus]
    [ 3.734802] platform_probe+0x3d/0x90
    [ 3.735684] really_probe+0x19b/0x400
    [ 3.736570] ? __device_attach_driver+0x100/0x100
    [ 3.737697] __driver_probe_device+0x78/0x160
```

```
     [ 3.738746] driver_probe_device+0x1f/0x90
     [ 3.739743] __driver_attach+0xc2/0x1b0
     [ 3.740671] bus_for_each_dev+0x70/0xc0
     [ 3.741601] bus_add_driver+0x10e/0x210
     [ 3.742527] driver_register+0x55/0xf0
     [ 3.744412] ? 0xffffffffc039a000
     [ 3.745207] hv_acpi_init+0x3c/0x1000 [hv_vmbus]

     Fixes: 7f163a6fd957 ("drivers:hv: Modify hv_vmbus to search for all MMIO ranges
 available.")
     Signed-off-by: Maciej S. Szmigiero <maciej.szmigiero@oracle.com>
     Reviewed-by: Michael Kelley <mikelley@microsoft.com>
     Signed-off-by: Wei Liu <wei.liu@kernel.org>
     Link:
 https://lore.kernel.org/r/fd8e64ceeecfd1d95ff49021080cf699e88dbbde.1691606267.git.maciej.szm
 igiero@oracle.com
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 44025cd15416c3cf6361a8292c3e6194b9022437
 Author: Arnd Bergmann <arnd@arndb.de>
 Date:   Mon Jul 24 21:56:49 2023 +0200

     remoteproc: stm32: fix incorrect optional pointers

     [ Upstream commit fb2bdd32b231b70e6a3f1054528692f604db25d8 ]

     Compile-testing without CONFIG_OF shows that the of_match_ptr() macro
     was used incorrectly here:

     drivers/remoteproc/stm32_rproc.c:662:34: warning: unused variable 'stm32_rproc_match' [-
 Wunused-const-variable]

     As in almost every driver, the solution is simply to remove the
     use of this macro. The same thing happened with the deprecated
     SIMPLE_DEV_PM_OPS(), but the corresponding warning was already shut
     up with __maybe_unused annotations, so fix those as well by using the
     correct DEFINE_SIMPLE_DEV_PM_OPS() macros and removing the extraneous
     __maybe_unused modifiers. For completeness, also add a pm_ptr() to let
     the PM ops be eliminated completely when CONFIG_PM is turned off.

     Reported-by: kernel test robot <lkp@intel.com>
     Closes: https://lore.kernel.org/oe-kbuild-all/202307242300.ia82qBTp-lkp@intel.com
     Fixes: 03bd158e1535 ("remoteproc: stm32: use correct format strings on 64-bit")
     Fixes: 410119ee29b6 ("remoteproc: stm32: wakeup the system by wdg irq")
     Fixes: 13140de09cc2 ("remoteproc: stm32: add an ST stm32_rproc driver")
     Signed-off-by: Arnd Bergmann <arnd@arndb.de>
     Acked-by: Arnaud Pouliquen <arnaud.pouliquen@foss.st.com>
     Link: https://lore.kernel.org/r/20230724195704.2432382-1-arnd@kernel.org
     Signed-off-by: Mathieu Poirier <mathieu.poirier@linaro.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 8430cb9b43d34f375239911457cd300da530af52
 Author: Fenghua Yu <fenghua.yu@intel.com>
 Date:   Thu Aug 10 18:26:35 2023 -0700

     dmaengine: idxd: Fix issues with PRS disable sysfs knob

     [ Upstream commit 8cae66574398326134a41513b419e00ad4e380ca ]

     There are two issues in the current PRS disable sysfs store function
     wq_prs_disable_store():

     1. Since PRS disable knob is invisible if PRS disable is not supported
        in WQ, it's redundant to check PRS support again in the store function
        again. Remove the redundant PRS support check.
```

```
    2. Since PRS disable is read-only when the device is not configurable,
       PRS disable cannot be changed on the device. Add device configurable
       check in the store function.

    Fixes: f2dc327131b5 ("dmaengine: idxd: add per wq PRS disable")
    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20230811012635.535413-2-fenghua.yu@intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b136a5b4892276eeca4d3d1ccb289f8b7c26d8c2
Author: Fenghua Yu <fenghua.yu@intel.com>
Date:   Thu Aug 10 18:26:34 2023 -0700

    dmaengine: idxd: Allow ATS disable update only for configurable devices

    [ Upstream commit 0056a7f07b0a63e6cee815a789eabba6f3a710f0 ]

    ATS disable status in a WQ is read-only if the device is not configurable.
    This change ensures that the ATS disable attribute can be modified via
    sysfs only on configurable devices.

    Fixes: 92de5fa2dc39 ("dmaengine: idxd: add ATS disable knob for work queues")
    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20230811012635.535413-1-fenghua.yu@intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c146070887950cec169639f53227b7cf0312ad19
Author: Fenghua Yu <fenghua.yu@intel.com>
Date:   Wed Jul 12 10:44:36 2023 -0700

    dmaengine: idxd: Expose ATS disable knob only when WQ ATS is supported

    [ Upstream commit 62b41b656666d2d35890124df5ef0881fe6d6769 ]

    WQ Advanced Translation Service (ATS) can be controlled only when
    WQ ATS is supported. The sysfs ATS disable knob should be visible only
    when the features is supported.

    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20230712174436.3435088-2-fenghua.yu@intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Stable-dep-of: 0056a7f07b0a ("dmaengine: idxd: Allow ATS disable update only for
configurable devices")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 733c6082c50a3007133315554bdeec58c740eabe
Author: Fenghua Yu <fenghua.yu@intel.com>
Date:   Wed Jul 12 10:44:35 2023 -0700

    dmaengine: idxd: Simplify WQ attribute visibility checks

    [ Upstream commit 97b1185fe54c8ce94104e3c7fa4ee0bbedd85920 ]

    The functions that check if WQ attributes are invisible are almost
    duplicate. Define a helper to simplify these functions and future
    WQ attribute visibility checks as well.

    Signed-off-by: Fenghua Yu <fenghua.yu@intel.com>
    Reviewed-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20230712174436.3435088-1-fenghua.yu@intel.com
```

```
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Stable-dep-of: 0056a7f07b0a ("dmaengine: idxd: Allow ATS disable update only for
 configurable devices")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 10826edc84eea8374dc7dba261f2291b023a616e
 Author: ruanjinjie <ruanjinjie@huawei.com>
 Date:   Mon Jul 24 14:41:08 2023 +0000

    dmaengine: ste_dma40: Add missing IRQ check in d40_probe

    [ Upstream commit c05ce6907b3d6e148b70f0bb5eafd61dcef1ddc1 ]

    Check for the return value of platform_get_irq(): if no interrupt
    is specified, it wouldn't make sense to call request_irq().

    Fixes: 8d318a50b3d7 ("DMAENGINE: Support for ST-Ericssons DMA40 block v3")
    Signed-off-by: Ruan Jinjie <ruanjinjie@huawei.com>
    Reviewed-by: Linus Walleij <linus.walleij@linaro.org>
    Link: https://lore.kernel.org/r/20230724144108.2582917-1-ruanjinjie@huawei.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 5fe497fb53eae9cec3f28d7a05fbee748ebc23ba
 Author: Randy Dunlap <rdunlap@infradead.org>
 Date:   Tue Aug 1 22:15:00 2023 -0700

    um: Fix hostaudio build errors

    [ Upstream commit db4bfcba7bb8d10f00bba2a3da6b9a9c2a1d7b71 ]

    Use "select" to ensure that the required kconfig symbols are set
    as expected.
    Drop HOSTAUDIO since it is now equivalent to UML_SOUND.

    Set CONFIG_SOUND=m in ARCH=um defconfig files to maintain the
    status quo of the default configs.

    Allow SOUND with UML regardless of HAS_IOMEM. Otherwise there is a
    kconfig warning for unmet dependencies. (This was not an issue when
    SOUND was defined in arch/um/drivers/Kconfig. I have done 50 randconfig
    builds and didn't find any issues.)

    This fixes build errors when CONFIG_SOUND is not set:

    ld: arch/um/drivers/hostaudio_kern.o: in function `hostaudio_cleanup_module':
    hostaudio_kern.c:(.exit.text+0xa): undefined reference to `unregister_sound_mixer'
    ld: hostaudio_kern.c:(.exit.text+0x15): undefined reference to `unregister_sound_dsp'
    ld: arch/um/drivers/hostaudio_kern.o: in function `hostaudio_init_module':
    hostaudio_kern.c:(.init.text+0x19): undefined reference to `register_sound_dsp'
    ld: hostaudio_kern.c:(.init.text+0x31): undefined reference to `register_sound_mixer'
    ld: hostaudio_kern.c:(.init.text+0x49): undefined reference to `unregister_sound_dsp'

    and this kconfig warning:
    WARNING: unmet direct dependencies detected for SOUND

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Fixes: d886e87cb82b ("sound: make OSS sound core optional")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: lore.kernel.org/r/202307141416.vxuRVpFv-lkp@intel.com
    Cc: Richard Weinberger <richard@nod.at>
    Cc: Anton Ivanov <anton.ivanov@cambridgegreys.com>
    Cc: Johannes Berg <johannes@sipsolutions.net>
    Cc: linux-um@lists.infradead.org
```

```
    Cc: Tejun Heo <tj@kernel.org>
    Cc: Takashi Iwai <tiwai@suse.de>
    Cc: Jaroslav Kysela <perex@perex.cz>
    Cc: Masahiro Yamada <masahiroy@kernel.org>
    Cc: Nathan Chancellor <nathan@kernel.org>
    Cc: Nick Desaulniers <ndesaulniers@google.com>
    Cc: Nicolas Schier <nicolas@fjasle.eu>
    Cc: linux-kbuild@vger.kernel.org
    Cc: alsa-devel@alsa-project.org
    Reviewed-by: Masahiro Yamada <masahiroy@kernel.org>
    Signed-off-by: Richard Weinberger <richard@nod.at>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c27237437b95b5afcf89d84fdc44b04ceecfe0a2
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Wed Jul 19 11:02:23 2023 +0200

    mfd: rz-mtu3: Link time dependencies

    [ Upstream commit 10d3340441bd0db857fc7fcb1733a800acf47a3d ]

    The new set of drivers for RZ/G2L MTU3a tries to enable compile-testing the
    individual client drivers even when the MFD portion is disabled but gets it
    wrong, causing a link failure when the core is in a loadable module but the
    other drivers are built-in:

    x86_64-linux-ld: drivers/pwm/pwm-rz-mtu3.o: in function `rz_mtu3_pwm_apply':
    pwm-rz-mtu3.c:(.text+0x4bf): undefined reference to `rz_mtu3_8bit_ch_write'
    x86_64-linux-ld: pwm-rz-mtu3.c:(.text+0x509): undefined reference to `rz_mtu3_disable'

    arm-linux-gnueabi-ld: drivers/counter/rz-mtu3-cnt.o: in function
`rz_mtu3_cascade_counts_enable_get':
    rz-mtu3-cnt.c:(.text+0xbec): undefined reference to `rz_mtu3_shared_reg_read'

    It seems better not to add the extra complexity here but instead just use
    a normal hard dependency, so remove the #else portion in the header along
    with the "|| COMPILE_TEST". This could also be fixed by having slightly more
    elaborate Kconfig dependencies or using the cursed 'IS_REACHABLE()' helper,
    but in practice it's already possible to compile-test all these drivers
    by enabling the mtd portion.

    Fixes: 254d3a727421c ("pwm: Add Renesas RZ/G2L MTU3a PWM driver")
    Fixes: 0be8907359df4 ("counter: Add Renesas RZ/G2L MTU3a counter driver")
    Fixes: 654c293e1687b ("mfd: Add Renesas RZ/G2L MTU3a core driver")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Acked-by: Thierry Reding <thierry.reding@gmail.com>
    Reviewed-by: Biju Das <biju.das.jz@bp.renesas.com>
    Link: https://lore.kernel.org/r/20230719090430.1925182-1-arnd@kernel.org
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8fceff5ef1fb64ec89b1e5a61f8e49c594084405
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Tue Jul 4 15:07:48 2023 +0200

    mfd: rk808: Make MFD_RK8XX tristate

    [ Upstream commit d085c27aa62999e2fe054707ab9da2af65d22b2f ]

    There is no reason for MFD_RK8XX to be bool, all drivers that depend on
    it, or that select it, are tristate.

    Fixes: c20e8c5b1203af37 ("mfd: rk808: Split into core and i2c")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Sebastian Reichel <sebastian.reichel@collabora.com>
```

```
    Tested-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Link:
https://lore.kernel.org/r/d132363fc9228473e9e652b70a3761b94de32d70.1688475844.git.geert+rene
sas@glider.be
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 527b1b3c8496dce12400a5f32eb41e70827c7995
Author: Yi Yang <yiyang13@huawei.com>
Date:   Thu Aug 17 19:58:39 2023 +0800

    mtd: rawnand: fsmc: handle clk prepare error in fsmc_nand_resume()

    [ Upstream commit a5a88125d00612586e941ae13e7fcf36ba8f18a7 ]

    In fsmc_nand_resume(), the return value of clk_prepare_enable() should be
    checked since it might fail.

    Fixes: e25da1c07dfb ("mtd: fsmc_nand: Add clk_{un}prepare() support")
    Signed-off-by: Yi Yang <yiyang13@huawei.com>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20230817115839.10192-1-yiyang13@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a0efa795fe9cff8a7b04185776b5ccfec6b547d5
Author: Hsin-Yi Wang <hsinyi@chromium.org>
Date:   Fri Aug 18 14:42:23 2023 +0800

    mtd: spi-nor: Check bus width while setting QE bit

    [ Upstream commit f01d8155a92e33cdaa85d20bfbe6c441907b3c1f ]

    spi_nor_write_16bit_sr_and_check() should also check if bus width is
    4 before setting QE bit.

    Fixes: 39d1e3340c73 ("mtd: spi-nor: Fix clearing of QE bit on lock()/unlock()")
    Suggested-by: Michael Walle <michael@walle.cc>
    Suggested-by: Tudor Ambarus <tudor.ambarus@linaro.org>
    Signed-off-by: Hsin-Yi Wang <hsinyi@chromium.org>
    Reviewed-by: Michael Walle <michael@walle.cc>
    Link: https://lore.kernel.org/r/20230818064524.1229100-2-hsinyi@chromium.org
    Signed-off-by: Tudor Ambarus <tudor.ambarus@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 84a283105483fb607884de50279fc225ff76e8a0
Author: Marek Behún <kabel@kernel.org>
Date:   Wed Aug 2 11:07:53 2023 +0200

    leds: trigger: tty: Do not use LED_ON/OFF constants, use led_blink_set_oneshot instead

    [ Upstream commit 730094577e0c37e1bc40be37cbd41f71b0a8a2a4 ]

    The tty LED trigger uses the obsolete LED_ON & LED_OFF constants when
    setting LED brightness. This is bad because the LED_ON constant is equal
    to 1, and so when activating the tty LED trigger on a LED class device
    with max_brightness greater than 1, the LED is dimmer than it can be
    (when max_brightness is 255, the LED is very dimm indeed; some devices
    translate 1/255 to 0, so the LED is OFF all the time).

    Instead of directly setting brightness to a specific value, use the
    led_blink_set_oneshot() function from LED core to configure the blink.
    This function takes the current configured brightness as blink
    brightness if not zero, and max brightness otherwise.

    This also changes the behavior of the TTY LED trigger. Previously if
```

```
    rx/tx stats kept changing, the LED was ON all the time they kept
    changing. With this patch the LED will blink on TTY activity.

    Fixes: fd4a641ac88f ("leds: trigger: implement a tty trigger")
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Link: https://lore.kernel.org/r/20230802090753.13611-1-kabel@kernel.org
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7db37d6583d0d6efff16366a06769a380a54a42d
Author: Marek Behún <kabel@kernel.org>
Date:   Tue Aug 1 17:16:23 2023 +0200

    leds: Fix BUG_ON check for LED_COLOR_ID_MULTI that is always false

    [ Upstream commit c3f853184bed04105682383c2971798c572226b5 ]

    At the time we call
        BUG_ON(props.color == LED_COLOR_ID_MULTI);
    the props variable is still initialized to zero.

    Call the BUG_ON only after we parse fwnode into props.

    Fixes: 77dce3a22e89 ("leds: disallow /sys/class/leds/*:multi:* for now")
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Link: https://lore.kernel.org/r/20230801151623.30387-1-kabel@kernel.org
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be3568956ed7b82e1e64cbd84b45665b3c998da7
Author: Marek Behún <kabel@kernel.org>
Date:   Tue Aug 1 14:49:31 2023 +0200

    leds: multicolor: Use rounded division when calculating color components

    [ Upstream commit 065d099f1be58187e6629273c50b948a02b7e1bf ]

    Given channel intensity, LED brightness and max LED brightness, the
    multicolor LED framework helper led_mc_calc_color_components() computes
    the color channel brightness as

        chan_brightness = brightness * chan_intensity / max_brightness

    Consider the situation when (brightness, intensity, max_brightness) is
    for example (16, 15, 255), then chan_brightness is computed to 0
    although the fractional divison would give 0.94, which should be rounded
    to 1.

    Use DIV_ROUND_CLOSEST here for the division to give more realistic
    component computation:

        chan_brightness = DIV_ROUND_CLOSEST(brightness * chan_intensity,
                                            max_brightness)

    Fixes: 55d5d3b46b08 ("leds: multicolor: Introduce a multicolor class definition")
    Signed-off-by: Marek Behún <kabel@kernel.org>
    Link: https://lore.kernel.org/r/20230801124931.8661-1-kabel@kernel.org
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc9be7140179595baaf7bb7ebc24dc368a714891
Author: Ahmad Fatoum <a.fatoum@pengutronix.de>
Date:   Sat Jul 8 13:26:46 2023 +0200

    thermal/drivers/imx8mm: Suppress log message on probe deferral
```

```
    [ Upstream commit 4afcb58ea47e66c025d2b0a5f091dce5aaf95b0f ]

    nvmem_cell_read_u32() may return -EPROBE_DEFER if NVMEM supplier has not
    yet been probed. Future reprobe may succeed, so printing:

      i.mx8mm_thermal 30260000.tmu: Failed to read OCOTP nvmem cell (-517).

    to the log is confusing. Fix this by using dev_err_probe. This also
    elevates the message from warning to error, which is more correct: The
    log message is only ever printed in probe error path and probe aborts
    afterwards, so it really warrants an error-level message.

    Fixes: 403291648823 ("thermal/drivers/imx: Add support for loading calibration data from
OCOTP")
    Signed-off-by: Ahmad Fatoum <a.fatoum@pengutronix.de>
    Reviewed-by: Marek Vasut <marex@denx.de>
    Reviewed-by: Peng Fan <peng.fan@nxp.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230708112647.2897294-1-a.fatoum@pengutronix.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ece35fb1ad12d1c5aea590fa5dfc9dde00454c7d
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:37 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Manage threshold between sensors

    [ Upstream commit 2bba1acf7a4cbe62abbb4c686e0414209ec5943b ]

    Each LVTS thermal controller can have up to four sensors, each capable
    of triggering its own interrupt when its measured temperature crosses
    the configured threshold. The threshold for each sensor is handled
    separately by the thermal framework, since each one is registered with
    its own thermal zone and trips. However, the temperature thresholds are
    configured on the controller, and therefore are shared between all
    sensors on that controller.

    When the temperature measured by the sensors is different enough to
    cause the thermal framework to configure different thresholds for each
    one, interrupts start triggering on sensors outside the last threshold
    configured.

    To address the issue, track the thresholds required by each sensor and
    only actually set the highest one in the hardware, and disable
    interrupts for all sensors outside the current configured range.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-7-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fd0f0f8f2b23c13298537cac72ac2c184cb7383a
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:36 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Don't leave threshold zeroed

    [ Upstream commit 77354eaef8218bc40d6b37e783b0b8dcca22a7d9 ]

    The thermal framework might leave the low threshold unset if there
```

```
    aren't any lower trip points. This leaves the register zeroed, which
    translates to a very high temperature for the low threshold. The
    interrupt for this threshold is then immediately triggered, and the
    state machine gets stuck, preventing any other temperature monitoring
    interrupts to ever trigger.

    (The same happens by not setting the Cold or Hot to Normal thresholds
    when using those)

    Set the unused threshold to a valid low value. This value was chosen so
    that for any valid golden temperature read from the efuse, when the
    value is converted to raw and back again to milliCelsius, the result
    doesn't underflow.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-6-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1bb31cf97a43a9a433a3cffa2ee10bf0420c80a
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:35 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Disable undesired interrupts

    [ Upstream commit 487bf099e85b724c824f5fafaf93c6749c4d2120 ]

    Out of the many interrupts supported by the hardware, the only ones of
    interest to the driver currently are:
    * The temperature went over the high offset threshold, for any of the
      sensors
    * The temperature went below the low offset threshold, for any of the
      sensors
    * The temperature went over the stage3 threshold

    These are the only thresholds configured by the driver through the
    OFFSETH, OFFSETL, and PROTTC registers, respectively.

    The current interrupt mask in LVTS_MONINT_CONF, enables many more
    interrupts, including data ready on sensors for both filtered and
    immediate mode. These are not only not handled by the driver, but they
    are also triggered too often, causing unneeded overhead. Disable these
    unnecessary interrupts.

    The meaning of each bit can be seen in the comment describing
    LVTS_MONINTST in the IRQ handler.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-5-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8493771a912379db0c96c225c4b2d3c1abaaeb96
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:34 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Use offset threshold for IRQ
```

```
    [ Upstream commit f79e996c7ed27bb196facbcd1c69ee33631d7051 ]

    There are two kinds of temperature monitoring interrupts available:
    * High Offset, Low Offset
    * Hot, Hot to normal, Cold

    The code currently uses the hot/h2n/cold interrupts, however in a way
    that doesn't work: the cold threshold is left uninitialized, which
    prevents the other thresholds from ever triggering, and the h2n
    interrupt is used as the lower threshold, which prevents the hot
    interrupt from triggering again after the thresholds are updated by the
    thermal framework, since a hot interrupt can only trigger again after
    the hot to normal interrupt has been triggered.

    But better yet than addressing those issues, is to use the high/low
    offset interrupts instead. This way only two thresholds need to be
    managed, which have a simpler state machine, making them a better match
    to the thermal framework's high and low thresholds.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-4-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit febd7d306972708e30890fb99708d709a834e113
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:33 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Honor sensors in immediate mode

    [ Upstream commit 64de162e34e4cb2982a1d96e492f018026a61c1d ]

    Each controller can be configured to operate on immediate or filtered
    mode. On filtered mode, the sensors are enabled by setting the
    corresponding bits in MONCTL0, while on immediate mode, by setting
    MSRCTL1.

    Previously, the code would set MSRCTL1 for all four sensors when
    configured to immediate mode, but given that the controller might not
    have all four sensors connected, this would cause interrupts to trigger
    for non-existent sensors. Fix this by handling the MSRCTL1 register
    analogously to the MONCTL0: only enable the sensors that were declared.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Tested-by: Chen-Yu Tsai <wenst@chromium.org>
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-3-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f987bf756630d463308ac0dfbc46b7baf3ce08a5
Author: Nícolas F. R. A. Prado <nfraprado@collabora.com>
Date:   Thu Jul 6 11:37:32 2023 -0400

    thermal/drivers/mediatek/lvts_thermal: Handle IRQ on all controllers

    [ Upstream commit cbd8c5aae2a988bafd4586bea710eeddc30a82ce ]
```

```
    There is a single IRQ handler for each LVTS thermal domain, and it is
    supposed to check each of its underlying controllers for the origin of
    the interrupt and clear its status. However due to a typo, only the
    first controller was ever being handled, which resulted in the interrupt
    never being cleared when it happened on the other controllers. Add the
    missing index so interrupts are handled for all controllers.

    Fixes: f5f633b18234 ("thermal/drivers/mediatek: Add the Low Voltage Thermal Sensor
driver")
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Tested-by: Chen-Yu Tsai <wenst@chromium.org>
    Signed-off-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Signed-off-by: Daniel Lezcano <daniel.lezcano@linaro.org>
    Link: https://lore.kernel.org/r/20230706153823.201943-2-nfraprado@collabora.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 146e244a00228c1dfb8a826825bb3abd3b69f39d
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Tue Jul 11 09:13:34 2023 +0300

    leds: pwm: Fix error code in led_pwm_create_fwnode()

    [ Upstream commit cadb2de2a7fd9e955381307de3eddfcc386c208e ]

    Negative -EINVAL was intended, not positive EINVAL.  Fix it.

    Fixes: 95138e01275e ("leds: pwm: Make error handling more robust")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Reviewed-by: Andy Shevchenko <andy.shevchenko@gmail.com>
    Link: https://lore.kernel.org/r/a33b981a-b2c4-4dc2-b00a-626a090d2f11@moroto.mountain
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aff4c1807fbab3da32afd2f2f881e4f7920d6f74
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Fri Jun 23 17:22:29 2023 +0200

    leds: simatic-ipc-leds-gpio: Restore LEDS_CLASS dependency

    [ Upstream commit 66c5e98bbf7b7b2ba0a095ef25bf55c7230e846e ]

    A recent rework accidentally lost the dependency on LEDS_CLASS, which
    leads to a link error when LED support is disbled:

    x86_64-linux-ld: drivers/leds/simple/simatic-ipc-leds.o: in function
`simatic_ipc_leds_probe':
    simatic-ipc-leds.c:(.text+0x10c): undefined reference to
`devm_led_classdev_register_ext'

    Add back the dependency that was there originally.

    Fixes: a6c80bec3c935 ("leds: simatic-ipc-leds-gpio: Add GPIO version of Siemens driver")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20230623152233.2246285-1-arnd@kernel.org
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a2f2e76ce2506d4547c2e58e7c34c474160626b
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Fri Jun 23 15:09:40 2023 +0300

    leds: aw200xx: Fix error code in probe()
```

```
    [ Upstream commit ad5152b85e8bc7dacb1e6e237553fbe779c938e0 ]

    The "ret" variable is zero/success here.  Don't return that, return
    -EINVAL instead.

    Fixes: 36a87f371b7a ("leds: Add AW20xx driver")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Link: https://lore.kernel.org/r/4d791b69-01c7-4532-818c-63712d3f63e1@moroto.mountain
    Signed-off-by: Lee Jones <lee@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3734a9558afac91df3c655a6f2376b9d14933b7
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Mon Jun 19 11:06:31 2023 +0800

    rpmsg: glink: Add check for kstrdup

    [ Upstream commit b5c9ee8296a3760760c7b5d2e305f91412adc795 ]

    Add check for the return value of kstrdup() and return the error
    if it fails in order to avoid NULL pointer dereference.

    Fixes: b4f8e52b89f6 ("rpmsg: Introduce Qualcomm RPM glink driver")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Link: https://lore.kernel.org/r/20230619030631.12361-1-jiasheng@iscas.ac.cn
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 18bd9f0a6fbde08d908ee4503186f763597c6524
Author: Jonas Karlman <jonas@kwiboo.se>
Date:   Thu Jun 15 17:10:21 2023 +0000

    phy/rockchip: inno-hdmi: do not power on rk3328 post pll on reg write

    [ Upstream commit 19a1d46bd699940a496d3b0d4e142ef99834988c ]

    inno_write is used to configure 0xaa reg, that also hold the
    POST_PLL_POWER_DOWN bit.
    When POST_PLL_REFCLK_SEL_TMDS is configured the power down bit is not
    taken into consideration.

    Fix this by keeping the power down bit until configuration is complete.
    Also reorder the reg write order for consistency.

    Fixes: 53706a116863 ("phy: add Rockchip Innosilicon hdmi phy")
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Link: https://lore.kernel.org/r/20230615171005.2251032-5-jonas@kwiboo.se
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a346f0ba56e3d828f1ef0bdf671c805c44ec9a3a
Author: Zheng Yang <zhengyang@rock-chips.com>
Date:   Thu Jun 15 17:10:19 2023 +0000

    phy/rockchip: inno-hdmi: round fractal pixclock in rk3328 recalc_rate

    [ Upstream commit d5ef343c1d62bc4c4c2c393af654a41cb34b449f ]

    inno_hdmi_phy_rk3328_clk_recalc_rate() is returning a rate not found
    in the pre pll config table when the fractal divider is used.
    This can prevent proper power_on because a tmdsclock for the new rate
    is not found in the pre pll config table.

    Fix this by saving and returning a rounded pixel rate that exist
```

```
    in the pre pll config table.

    Fixes: 53706a116863 ("phy: add Rockchip Innosilicon hdmi phy")
    Signed-off-by: Zheng Yang <zhengyang@rock-chips.com>
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Link: https://lore.kernel.org/r/20230615171005.2251032-3-jonas@kwiboo.se
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 347edd711b2b65313759edfb827139908f3e0098
Author: Jonas Karlman <jonas@kwiboo.se>
Date:   Thu Jun 15 17:10:17 2023 +0000

    phy/rockchip: inno-hdmi: use correct vco_div_5 macro on rk3328

    [ Upstream commit 644c06dfbd0da713f772abf0a8f8581ac78e6264 ]

    inno_hdmi_phy_rk3328_clk_set_rate() is using the RK3228 macro
    when configuring vco_div_5 on RK3328.

    Fix this by using correct vco_div_5 macro for RK3328.

    Fixes: 53706a116863 ("phy: add Rockchip Innosilicon hdmi phy")
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Link: https://lore.kernel.org/r/20230615171005.2251032-2-jonas@kwiboo.se
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1340b98fa7f7396c183460f2e972bf4c8100cf0d
Author: Rex Zhang <rex.zhang@intel.com>
Date:   Wed Jun 14 14:27:06 2023 +0800

    dmaengine: idxd: Modify the dependence of attribute pasid_enabled

    [ Upstream commit 50c5e6f41d5ad7c731c31135a30d0e4f0e4fea26 ]

    Kernel PASID and user PASID are separately enabled. User needs to know the
    user PASID enabling status to decide how to use IDXD device in user space.
    This is done via the attribute /sys/bus/dsa/devices/dsa0/pasid_enabled.
    It's unnecessary for user to know the kernel PASID enabling status because
    user won't use the kernel PASID. But instead of showing the user PASID
    enabling status, the attribute shows the kernel PASID enabling status. Fix
    the issue by showing the user PASID enabling status in the attribute.

    Fixes: 42a1b73852c4 ("dmaengine: idxd: Separate user and kernel pasid enabling")
    Signed-off-by: Rex Zhang <rex.zhang@intel.com>
    Acked-by: Fenghua Yu <fenghua.yu@intel.com>
    Acked-by: Dave Jiang <dave.jiang@intel.com>
    Link: https://lore.kernel.org/r/20230614062706.1743078-1-rex.zhang@intel.com
    Signed-off-by: Vinod Koul <vkoul@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4977b81fc58d5b6d790f9a94d26fa01bdf112f5e
Author: William Zhang <william.zhang@broadcom.com>
Date:   Thu Jul 6 11:29:09 2023 -0700

    mtd: rawnand: brcmnand: Fix mtd oobsize

    [ Upstream commit 60177390fa061c62d156f4a546e3efd90df3c183 ]

    brcmnand controller can only access the flash spare area up to certain
    bytes based on the ECC level. It can be less than the actual flash spare
    area size. For example, for many NAND chip supporting ECC BCH-8, it has
    226 bytes spare area. But controller can only uses 218 bytes. So brcmand
    driver overrides the mtd oobsize with the controller's accessible spare
```

```
    area size. When the nand base driver utilizes the nand_device object, it
    resets the oobsize back to the actual flash spare aprea size from
    nand_memory_organization structure and controller may not able to access
    all the oob area as mtd advises.

    This change fixes the issue by overriding the oobsize in the
    nand_memory_organization structure to the controller's accessible spare
    area size.

    Fixes: a7ab085d7c16 ("mtd: rawnand: Initialize the nand_device object")
    Signed-off-by: William Zhang <william.zhang@broadcom.com>
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Link: https://lore.kernel.org/linux-mtd/20230706182909.79151-6-
william.zhang@broadcom.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 24e84a96863aaa9809a9a8919921e0b94ac6a51a
Author: Hugh Dickins <hughd@google.com>
Date:   Sat Sep 2 08:29:30 2023 -0700

    mm/pagewalk: fix bootstopping regression from extra pte_unmap()

    [ Upstream commit ee40d543e97d23d3392d8fb1ec9972eb4e9c7611 ]

    Mikhail reports early-6.6-based Fedora Rawhide not booting: "rcu_preempt
    detected expedited stalls", minutes wait, and then hung_task splat while
    kworker trying to synchronize_rcu_expedited().  Nothing logged to disk.

    He bisected to my 6.6 a349d72fd9ef ("mm/pgtable: add rcu_read_lock() and
    rcu_read_unlock()s"): but the one to blame is my 6.5 commit to fix the
    espfix "bad pmd" warnings when booting x86_64 with CONFIG_EFI_PGT_DUMP=y.

    Gaah, that added an "addr >= TASK_SIZE" check to avoid pte_offset_map(),
    but failed to add the equivalent check when choosing to pte_unmap().

    It's not a problem on 6.5 (for different reasons, it's harmless on both
    64-bit and 32-bit), but becomes a bootstopper on 6.6 with the unbalanced
    rcu_read_unlock() - RCU has a WARN_ON_ONCE for that, but it would have
    scrolled off Mikhail's console too quickly.

    Reported-by: Mikhail Gavrilov <mikhail.v.gavrilov@gmail.com>
    Closes: https://lore.kernel.org/linux-
mm/CABXGCsNi8Tiv5zUPNXr6UJw6qV1VdaBEfGqEAMkkXE3QPvZuAQ@mail.gmail.com/
    Fixes: 8b1cb4a2e819 ("mm/pagewalk: fix EFI_PGT_DUMP of espfix area")
    Fixes: a349d72fd9ef ("mm/pgtable: add rcu_read_lock() and rcu_read_unlock()s")
    Signed-off-by: Hugh Dickins <hughd@google.com>
    Tested-by: Mikhail Gavrilov <mikhail.v.gavrilov@gmail.com>
    Signed-off-by: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37ca1b686078b00cc4ffa008e2190615f7709b5d
Author: Zheng Yejian <zhengyejian1@huawei.com>
Date:   Thu Aug 31 21:27:39 2023 +0800

    tracing: Fix race issue between cpu buffer write and swap

    [ Upstream commit 3163f635b20e9e1fb4659e74f47918c9dddfe64e ]

    Warning happened in rb_end_commit() at code:
            if (RB_WARN_ON(cpu_buffer, !local_read(&cpu_buffer->committing)))

      WARNING: CPU: 0 PID: 139 at kernel/trace/ring_buffer.c:3142
            rb_commit+0x402/0x4a0
      Call Trace:
       ring_buffer_unlock_commit+0x42/0x250
```

```
 trace_buffer_unlock_commit_regs+0x3b/0x250
 trace_event_buffer_commit+0xe5/0x440
 trace_event_buffer_reserve+0x11c/0x150
 trace_event_raw_event_sched_switch+0x23c/0x2c0
 __traceiter_sched_switch+0x59/0x80
 __schedule+0x72b/0x1580
 schedule+0x92/0x120
 worker_thread+0xa0/0x6f0
```

It is because the race between writing event into cpu buffer and swapping cpu buffer through file per_cpu/cpu0/snapshot:

```
  Write on CPU 0             Swap buffer by per_cpu/cpu0/snapshot on CPU 1
  --------                   --------
                             tracing_snapshot_write()
                               [...]

  ring_buffer_lock_reserve()
    cpu_buffer = buffer->buffers[cpu]; // 1. Suppose find 'cpu_buffer_a';
    [...]
    rb_reserve_next_event()
      [...]

                               ring_buffer_swap_cpu()
                                 if (local_read(&cpu_buffer_a->committing))
                                     goto out_dec;
                                 if (local_read(&cpu_buffer_b->committing))
                                     goto out_dec;
                                 buffer_a->buffers[cpu] = cpu_buffer_b;
                                 buffer_b->buffers[cpu] = cpu_buffer_a;
                                 // 2. cpu_buffer has swapped here.

      rb_start_commit(cpu_buffer);
      if (unlikely(READ_ONCE(cpu_buffer->buffer)
          != buffer)) { // 3. This check passed due to 'cpu_buffer->buffer'
        [...]           //    has not changed here.
        return NULL;
      }
                                 cpu_buffer_b->buffer = buffer_a;
                                 cpu_buffer_a->buffer = buffer_b;
                                 [...]

      // 4. Reserve event from 'cpu_buffer_a'.

  ring_buffer_unlock_commit()
    [...]
    cpu_buffer = buffer->buffers[cpu]; // 5. Now find 'cpu_buffer_b' !!!
    rb_commit(cpu_buffer)
      rb_end_commit()  // 6. WARN for the wrong 'committing' state !!!
```

Based on above analysis, we can easily reproduce by following testcase:

```
  ``` bash
  #!/bin/bash

  dmesg -n 7
  sysctl -w kernel.panic_on_warn=1
  TR=/sys/kernel/tracing
  echo 7 > ${TR}/buffer_size_kb
  echo "sched:sched_switch" > ${TR}/set_event
  while [ true ]; do
          echo 1 > ${TR}/per_cpu/cpu0/snapshot
  done &
  while [ true ]; do
          echo 1 > ${TR}/per_cpu/cpu0/snapshot
  done &
```
```

```
       while [ true ]; do
               echo 1 > ${TR}/per_cpu/cpu0/snapshot
       done &
       ```

    To fix it, IIUC, we can use smp_call_function_single() to do the swap on
    the target cpu where the buffer is located, so that above race would be
    avoided.

    Link: https://lore.kernel.org/linux-trace-kernel/20230831132739.4070878-1-
zhengyejian1@huawei.com

    Cc: <mhiramat@kernel.org>
    Fixes: f1affcaaa861 ("tracing: Add snapshot in the per_cpu trace directories")
    Signed-off-by: Zheng Yejian <zhengyejian1@huawei.com>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41b9a21be4401225f1156b0540be11db07138dab
Author: Mikhail Kobuk <m.kobuk@ispras.ru>
Date: Fri Aug 25 13:34:30 2023 +0300

    tracing: Remove extra space at the end of hwlat_detector/mode

    [Upstream commit 2cf0dee989a8b2501929eaab29473b6b1fa11057]

    Space is printed after each mode value including the last one:
    $ echo \"$(sudo cat /sys/kernel/tracing/hwlat_detector/mode)\"
    "none [round-robin] per-cpu "

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Link: https://lore.kernel.org/linux-trace-kernel/20230825103432.7750-1-m.kobuk@ispras.ru

    Cc: Masami Hiramatsu <mhiramat@kernel.org>
    Fixes: 8fa826b7344d ("trace/hwlat: Implement the mode config option")
    Signed-off-by: Mikhail Kobuk <m.kobuk@ispras.ru>
    Reviewed-by: Alexey Khoroshilov <khoroshilov@ispras.ru>
    Acked-by: Daniel Bristot de Oliveira <bristot@kernel.org>
    Signed-off-by: Steven Rostedt (Google) <rostedt@goodmis.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6a734ca06b108fc0311278e7c08ade74de9db035
Author: Dave Hansen <dave.hansen@linux.intel.com>
Date: Tue Aug 29 08:07:25 2023 -0700

    x86/speculation: Mark all Skylake CPUs as vulnerable to GDS

    [Upstream commit c9f4c45c8ec3f07f4f083f9750032a1ec3eab6b2]

    The Gather Data Sampling (GDS) vulnerability is common to all Skylake
    processors. However, the "client" Skylakes* are now in this list:

https://www.intel.com/content/www/us/en/support/articles/000022396/processors.html

    which means they are no longer included for new vulnerabilities here:

             https://www.intel.com/content/www/us/en/developer/topic-technology/software-
security-guidance/processors-affected-consolidated-product-cpu-model.html

    or in other GDS documentation. Thus, they were not included in the
    original GDS mitigation patches.

    Mark SKYLAKE and SKYLAKE_L as vulnerable to GDS to match all the


```
```

```
    other Skylake CPUs (which include Kaby Lake).  Also group the CPUs
    so that the ones that share the exact same vulnerabilities are next
    to each other.

    Last, move SRBDS to the end of each line.  This makes it clear at a
    glance that SKYLAKE_X is unique.  Of the five Skylakes, it is the
    only "server" CPU and has a different implementation from the
    clients of the "special register" hardware, making it immune to SRBDS.

    This makes the diff much harder to read, but the resulting table is
    worth it.

    I very much appreciate the report from Michael Zhivich about this
    issue.  Despite what level of support a hardware vendor is providing,
    the kernel very much needs an accurate and up-to-date list of
    vulnerable CPUs.  More reports like this are very welcome.

    * Client Skylakes are CPUID 406E3/506E3 which is family 6, models
      0x4E and 0x5E, aka INTEL_FAM6_SKYLAKE and INTEL_FAM6_SKYLAKE_L.

    Reported-by: Michael Zhivich <mzhivich@akamai.com>
    Fixes: 8974eb588283 ("x86/speculation: Add Gather Data Sampling mitigation")
    Signed-off-by: Dave Hansen <dave.hansen@linux.intel.com>
    Signed-off-by: Ingo Molnar <mingo@kernel.org>
    Reviewed-by: Daniel Sneddon <daniel.sneddon@linux.intel.com>
    Cc: Linus Torvalds <torvalds@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f89944f2aecfbb8ef23ec467a7d43c8737959477
Author: Paul Gortmaker <paul.gortmaker@windriver.com>
Date:   Fri Aug 18 16:07:57 2023 -0400

    tick/rcu: Fix false positive "softirq work is pending" messages

    [ Upstream commit 96c1fa04f089a7e977a44e4e8fdc92e81be20bef ]

    In commit 0345691b24c0 ("tick/rcu: Stop allowing RCU_SOFTIRQ in idle") the
    new function report_idle_softirq() was created by breaking code out of the
    existing can_stop_idle_tick() for kernels v5.18 and newer.

    In doing so, the code essentially went from a one conditional:

            if (a && b && c)
                    warn();

    to a three conditional:

            if (!a)
                    return;
            if (!b)
                    return;
            if (!c)
                    return;
            warn();

    But that conversion got the condition for the RT specific
    local_bh_blocked() wrong. The original condition was:

            !local_bh_blocked()

    but the conversion failed to negate it so it ended up as:

            if (!local_bh_blocked())
                    return false;
```

```
    This issue lay dormant until another fixup for the same commit was added
    in commit a7e282c77785 ("tick/rcu: Fix bogus ratelimit condition").
    This commit realized the ratelimit was essentially set to zero instead
    of ten, and hence *no* softirq pending messages would ever be issued.

    Once this commit was backported via linux-stable, both the v6.1 and v6.4
    preempt-rt kernels started printing out 10 instances of this at boot:

      NOHZ tick-stop error: local softirq work is pending, handler #80!!!

    Remove the negation and return when local_bh_blocked() evaluates to true to
    bring the correct behaviour back.

    Fixes: 0345691b24c0 ("tick/rcu: Stop allowing RCU_SOFTIRQ in idle")
    Signed-off-by: Paul Gortmaker <paul.gortmaker@windriver.com>
    Signed-off-by: Thomas Gleixner <tglx@linutronix.de>
    Tested-by: Ahmad Fatoum <a.fatoum@pengutronix.de>
    Reviewed-by: Wen Yang <wenyang.linux@foxmail.com>
    Acked-by: Frederic Weisbecker <frederic@kernel.org>
    Link: https://lore.kernel.org/r/20230818200757.1808398-1-paul.gortmaker@windriver.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ce55024f28589b0012fa2c6b5748ec5a180b7fbe
Author: Mirsad Goran Todorovac <mirsad.todorovac@alu.unizg.hr>
Date:   Sat Aug 26 16:51:03 2023 +0200

    workqueue: fix data race with the pwq->stats[] increment

    [ Upstream commit fe48ba7daefe75bbbefa2426deddc05f2d530d2d ]

    KCSAN has discovered a data race in kernel/workqueue.c:2598:

    [ 1863.554079] ==================================================================
    [ 1863.554118] BUG: KCSAN: data-race in process_one_work / process_one_work

    [ 1863.554142] write to 0xffff963d99d79998 of 8 bytes by task 5394 on cpu 27:
    [ 1863.554154] process_one_work (kernel/workqueue.c:2598)
    [ 1863.554166] worker_thread (./include/linux/list.h:292 kernel/workqueue.c:2752)
    [ 1863.554177] kthread (kernel/kthread.c:389)
    [ 1863.554186] ret_from_fork (arch/x86/kernel/process.c:145)
    [ 1863.554197] ret_from_fork_asm (arch/x86/entry/entry_64.S:312)

    [ 1863.554213] read to 0xffff963d99d79998 of 8 bytes by task 5450 on cpu 12:
    [ 1863.554224] process_one_work (kernel/workqueue.c:2598)
    [ 1863.554235] worker_thread (./include/linux/list.h:292 kernel/workqueue.c:2752)
    [ 1863.554247] kthread (kernel/kthread.c:389)
    [ 1863.554255] ret_from_fork (arch/x86/kernel/process.c:145)
    [ 1863.554266] ret_from_fork_asm (arch/x86/entry/entry_64.S:312)

    [ 1863.554280] value changed: 0x0000000000001766 -> 0x000000000000176a

    [ 1863.554295] Reported by Kernel Concurrency Sanitizer on:
    [ 1863.554303] CPU: 12 PID: 5450 Comm: kworker/u64:1 Tainted: G             L     6.5.0-
rc6+ #44
    [ 1863.554314] Hardware name: ASRock X670E PG Lightning/X670E PG Lightning, BIOS 1.21
04/26/2023
    [ 1863.554322] Workqueue: btrfs-endio btrfs_end_bio_work [btrfs]
    [ 1863.554941] ==================================================================

        lockdep_invariant_state(true);
    →   pwq->stats[PWQ_STAT_STARTED]++;
        trace_workqueue_execute_start(work);
        worker->current_func(work);

    Moving pwq->stats[PWQ_STAT_STARTED]++; before the line
```

```
         raw_spin_unlock_irq(&pool->lock);
```

    resolves the data race without performance penalty.

    KCSAN detected at least one additional data race:

```
    [  157.834751] ==================================================================
    [  157.834770] BUG: KCSAN: data-race in process_one_work / process_one_work

    [  157.834793] write to 0xffff9934453f77a0 of 8 bytes by task 468 on cpu 29:
    [  157.834804] process_one_work
 (/home/marvin/linux/kernel/linux_torvalds/kernel/workqueue.c:2606)
    [  157.834815] worker_thread
 (/home/marvin/linux/kernel/linux_torvalds/./include/linux/list.h:292
 /home/marvin/linux/kernel/linux_torvalds/kernel/workqueue.c:2752)
    [  157.834826] kthread (/home/marvin/linux/kernel/linux_torvalds/kernel/kthread.c:389)
    [  157.834834] ret_from_fork
 (/home/marvin/linux/kernel/linux_torvalds/arch/x86/kernel/process.c:145)
    [  157.834845] ret_from_fork_asm
 (/home/marvin/linux/kernel/linux_torvalds/arch/x86/entry/entry_64.S:312)

    [  157.834859] read to 0xffff9934453f77a0 of 8 bytes by task 214 on cpu 7:
    [  157.834868] process_one_work
 (/home/marvin/linux/kernel/linux_torvalds/kernel/workqueue.c:2606)
    [  157.834879] worker_thread
 (/home/marvin/linux/kernel/linux_torvalds/./include/linux/list.h:292
 /home/marvin/linux/kernel/linux_torvalds/kernel/workqueue.c:2752)
    [  157.834890] kthread (/home/marvin/linux/kernel/linux_torvalds/kernel/kthread.c:389)
    [  157.834897] ret_from_fork
 (/home/marvin/linux/kernel/linux_torvalds/arch/x86/kernel/process.c:145)
    [  157.834907] ret_from_fork_asm
 (/home/marvin/linux/kernel/linux_torvalds/arch/x86/entry/entry_64.S:312)

    [  157.834920] value changed: 0x000000000000052a -> 0x0000000000000532

    [  157.834933] Reported by Kernel Concurrency Sanitizer on:
    [  157.834941] CPU: 7 PID: 214 Comm: kworker/u64:2 Tainted: G             L     6.5.0-
 rc7-kcsan-00169-g81eaf55a60fc #4
    [  157.834951] Hardware name: ASRock X670E PG Lightning/X670E PG Lightning, BIOS 1.21
 04/26/2023
    [  157.834958] Workqueue: btrfs-endio btrfs_end_bio_work [btrfs]
    [  157.835567] ==================================================================
```

    in code:

```
            trace_workqueue_execute_end(work, worker->current_func);
    →       pwq->stats[PWQ_STAT_COMPLETED]++;
            lock_map_release(&lockdep_map);
            lock_map_release(&pwq->wq->lockdep_map);
```

    which needs to be resolved separately.

    Fixes: 725e8ec59c56c ("workqueue: Add pwq->stats[] and a monitoring script")
    Cc: Tejun Heo <tj@kernel.org>
    Suggested-by: Lai Jiangshan <jiangshanlai@gmail.com>
    Link: https://lore.kernel.org/lkml/20230818194448.29672-1-mirsad.todorovac@alu.unizg.hr/
    Signed-off-by: Mirsad Goran Todorovac <mirsad.todorovac@alu.unizg.hr>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

```
commit 37a42eb48fbb5e96b70dd6afa5a251bbff902a61
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Wed Aug 23 13:54:21 2023 -0500
```

```
    platform/x86/amd/pmf: Fix a missing cleanup path

    [ Upstream commit 4dbd6e61adc7e52dd1c9165f0ccaa90806611e40 ]

    On systems that support slider notifications but don't otherwise support
    granular slider the SPS cleanup path doesn't run.

    This means that loading/unloading/loading leads to failures because
    the sysfs files don't get setup properly when reloaded.

    Add the missing cleanup path.

    Fixes: 33c9ab5b493a ("platform/x86/amd/pmf: Notify OS power slider update")
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Link: https://lore.kernel.org/r/20230823185421.23959-1-mario.limonciello@amd.com
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b85d3807e5ec368bfd5b20245347d7c1434aff76
Author: Rahul Rameshbabu <rrameshbabu@nvidia.com>
Date:   Thu Aug 24 06:14:54 2023 +0000

    HID: nvidia-shield: Reference hid_device devm allocation of input_dev name

    [ Upstream commit 197d3143520fec9fde89aebabc9f0d7464f08e50 ]

    Use hid_device for devm allocation of the input_dev name to avoid a
    use-after-free. input_unregister_device would trigger devres cleanup of all
    resources associated with the input_dev, free-ing the name. The name would
    subsequently be used in a uevent fired at the end of unregistering the
    input_dev.

    Reported-by: Maxime Ripard <mripard@kernel.org>
    Closes: https://lore.kernel.org/linux-
input/ZOZIZCND+L0P1wJc@penguin/T/#m443f3dce92520f74b6cf6ffa8653f9c92643d4ae
    Fixes: 09308562d4af ("HID: nvidia-shield: Initial driver implementation with
Thunderstrike support")
    Suggested-by: Maxime Ripard <mripard@kernel.org>
    Suggested-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Rahul Rameshbabu <rrameshbabu@nvidia.com>
    Reviewed-by: Maxime Ripard <mripard@kernel.org>
    Link: https://lore.kernel.org/r/20230824061308.222021-4-sergeantsagara@protonmail.com
    Signed-off-by: Benjamin Tissoires <bentiss@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b70ac7849248ec8128fa12f86e3655ba38838f29
Author: Rahul Rameshbabu <sergeantsagara@protonmail.com>
Date:   Thu Aug 24 06:14:33 2023 +0000

    HID: multitouch: Correct devm device reference for hidinput input_dev name

    [ Upstream commit 4794394635293a3e74591351fff469cea7ad15a2 ]

    Reference the HID device rather than the input device for the devm
    allocation of the input_dev name. Referencing the input_dev would lead to a
    use-after-free when the input_dev was unregistered and subsequently fires a
    uevent that depends on the name. At the point of firing the uevent, the
    name would be freed by devres management.

    Use devm_kasprintf to simplify the logic for allocating memory and
    formatting the input_dev name string.

    Reported-by: Maxime Ripard <mripard@kernel.org>
    Closes: https://lore.kernel.org/linux-
```

```
  input/ZOZIZCND+L0P1wJc@penguin/T/#m443f3dce92520f74b6cf6ffa8653f9c92643d4ae
      Fixes: c08d46aa805b ("HID: multitouch: devm conversion")
      Suggested-by: Maxime Ripard <mripard@kernel.org>
      Suggested-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
      Signed-off-by: Rahul Rameshbabu <sergeantsagara@protonmail.com>
      Reviewed-by: Maxime Ripard <mripard@kernel.org>
      Link: https://lore.kernel.org/r/20230824061308.222021-3-sergeantsagara@protonmail.com
      Signed-off-by: Benjamin Tissoires <bentiss@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 58f0d1c0e494a88f301bf455da7df4366f179bbb
  Author: Rahul Rameshbabu <sergeantsagara@protonmail.com>
  Date:   Thu Aug 24 06:14:17 2023 +0000

      HID: uclogic: Correct devm device reference for hidinput input_dev name

      [ Upstream commit dd613a4e45f8d35f49a63a2064e5308fa5619e29 ]

      Reference the HID device rather than the input device for the devm
      allocation of the input_dev name. Referencing the input_dev would lead to a
      use-after-free when the input_dev was unregistered and subsequently fires a
      uevent that depends on the name. At the point of firing the uevent, the
      name would be freed by devres management.

      Use devm_kasprintf to simplify the logic for allocating memory and
      formatting the input_dev name string.

      Reported-by: syzbot+3a0ebe8a52b89c63739d@syzkaller.appspotmail.com
      Closes: https://lore.kernel.org/linux-input/ZOZIZCND+L0P1wJc@penguin/T/
      Reported-by: Maxime Ripard <mripard@kernel.org>
      Closes: https://lore.kernel.org/linux-
  input/ZOZIZCND+L0P1wJc@penguin/T/#m443f3dce92520f74b6cf6ffa8653f9c92643d4ae
      Fixes: cce2dbdf258e ("HID: uclogic: name the input nodes based on their tool")
      Suggested-by: Maxime Ripard <mripard@kernel.org>
      Suggested-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
      Signed-off-by: Rahul Rameshbabu <sergeantsagara@protonmail.com>
      Reviewed-by: Maxime Ripard <mripard@kernel.org>
      Link: https://lore.kernel.org/r/20230824061308.222021-2-sergeantsagara@protonmail.com
      Signed-off-by: Benjamin Tissoires <bentiss@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 497dd09efc30182e55c258a0f08e1206317141b1
  Author: Chao Yu <chao@kernel.org>
  Date:   Mon Aug 21 23:22:25 2023 +0800

      f2fs: compress: fix to assign compress_level for lz4 correctly

      [ Upstream commit 091a4dfbb1d32b06c031edbfe2a44af100c4604f ]

      After remount, F2FS_OPTION().compress_level was assgin to
      LZ4HC_DEFAULT_CLEVEL incorrectly, result in lz4hc:9 was enabled, fix it.

      1. mount /dev/vdb
      /dev/vdb on /mnt/f2fs type f2fs (...,compress_algorithm=lz4,compress_log_size=2,...)
      2. mount -t f2fs -o remount,compress_log_size=3 /mnt/f2fs/
      3. mount|grep f2fs
      /dev/vdb on /mnt/f2fs type f2fs (...,compress_algorithm=lz4:9,compress_log_size=3,...)

      Fixes: 00e120b5e4b5 ("f2fs: assign default compression level")
      Signed-off-by: Chao Yu <chao@kernel.org>
      Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 825588e92ab01862e9464c134be7f1ba5b49db2e
  Author: Björn Töpel <bjorn@rivosinc.com>
```

```
Date:   Wed Aug 23 10:28:45 2023 +0200

    riscv: Require FRAME_POINTER for some configurations

    [ Upstream commit 9f944d2e0ab39296bfadb29167dc333815ba9f48 ]

    Some V configurations implicitly turn on '-fno-omit-frame-pointer',
    but leaving FRAME_POINTER disabled. This makes it hard to reason about
    the FRAME_POINTER config, and also triggers build failures introduced
    in by the commit in the Fixes: tag.

    Select FRAME_POINTER explicitly for these configurations.

    Fixes: ebc9cb03b21e ("riscv: stack: Fixup independent softirq stack for
CONFIG_FRAME_POINTER=n")
    Signed-off-by: Björn Töpel <bjorn@rivosinc.com>
    Tested-by: Randy Dunlap <rdunlap@infradead.org>
    Acked-by: Randy Dunlap <rdunlap@infradead.org>
    Reviewed-by: Conor Dooley <conor.dooley@microchip.com>
    Link: https://lore.kernel.org/r/20230823082845.354839-1-bjorn@kernel.org
    Signed-off-by: Palmer Dabbelt <palmer@rivosinc.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9529b1c5c38f481c74440b3f018ea3d20d5d3320
Author: Miquel Raynal <miquel.raynal@bootlin.com>
Date:   Wed Aug 23 14:27:42 2023 +0100

    nvmem: core: Return NULL when no nvmem layout is found

    [ Upstream commit 81e1d9a39569d315f747c2af19ce502cd08645ed ]

    Currently, of_nvmem_layout_get_container() returns NULL on error, or an
    error pointer if either CONFIG_NVMEM or CONFIG_OF is turned off. We
    should likely avoid this kind of mix for two reasons: to clarify the
    intend and anyway fix the !CONFIG_OF which will likely always if we use
    this helper somewhere else. Let's just return NULL when no layout is
    found, we don't need an error value here.

    Link: https://staticthinking.wordpress.com/2022/08/01/mixing-error-pointers-and-null/
    Fixes: 266570f496b9 ("nvmem: core: introduce NVMEM layouts")
    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <dan.carpenter@linaro.org>
    Closes: https://lore.kernel.org/r/202308030002.DnSFOrMB-lkp@intel.com/
    Signed-off-by: Miquel Raynal <miquel.raynal@bootlin.com>
    Reviewed-by: Michael Walle <michael@walle.cc>
    Signed-off-by: Srinivas Kandagatla <srinivas.kandagatla@linaro.org>
    Link: https://lore.kernel.org/r/20230823132744.350618-21-srinivas.kandagatla@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00ab92481d3a40a5ad323df4c518068f66ce49f1
Author: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
Date:   Tue Jun 13 03:16:35 2023 -0700

    HID: logitech-dj: Fix error handling in logi_dj_recv_switch_to_dj_mode()

    [ Upstream commit 6f20d3261265885f6a6be4cda49d7019728760e0 ]

    Presently, if a call to logi_dj_recv_send_report() fails, we do
    not learn about the error until after sending short
    HID_OUTPUT_REPORT with hid_hw_raw_request().
    To handle this somewhat unlikely issue, return on error in
    logi_dj_recv_send_report() (minding ugly sleep workaround) and
    take into account the result of hid_hw_raw_request().
```

```
    Found by Linux Verification Center (linuxtesting.org) with static
    analysis tool SVACE.

    Fixes: 6a9ddc897883 ("HID: logitech-dj: enable notifications on connect/disconnect")
    Signed-off-by: Nikita Zhandarovich <n.zhandarovich@fintech.ru>
    Link: https://lore.kernel.org/r/20230613101635.77820-1-n.zhandarovich@fintech.ru
    Signed-off-by: Benjamin Tissoires <bentiss@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 24f9884971f9b34915b67baacf7350a3f6f19ea4
Author: Yonatan Nachum <ynachum@amazon.com>
Date:   Tue Aug 22 08:27:25 2023 +0000

    RDMA/efa: Fix wrong resources deallocation order

    [ Upstream commit dc202c57e9a1423aed528e4b8dc949509cd32191 ]

    When trying to destroy QP or CQ, we first decrease the refcount and
    potentially free memory regions allocated for the object and then
    request the device to destroy the object. If the device fails, the
    object isn't fully destroyed so the user/IB core can try to destroy the
    object again which will lead to underflow when trying to decrease an
    already zeroed refcount.

    Deallocate resources in reverse order of allocating them to safely free
    them.

    Fixes: ff6629f88c52 ("RDMA/efa: Do not delay freeing of DMA pages")
    Reviewed-by: Michael Margolin <mrgolin@amazon.com>
    Reviewed-by: Yossi Leybovich <sleybo@amazon.com>
    Signed-off-by: Yonatan Nachum <ynachum@amazon.com>
    Link: https://lore.kernel.org/r/20230822082725.31719-1-ynachum@amazon.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 782f3b65f3ade065f12fe778d3bbda304f6686cb
Author: Guoqing Jiang <guoqing.jiang@linux.dev>
Date:   Mon Aug 21 21:32:54 2023 +0800

    RDMA/siw: Correct wrong debug message

    [ Upstream commit bee024d20451e4ce04ea30099cad09f7f75d288b ]

    We need to print num_sle first then pbl->max_buf per the condition.
    Also replace mem->pbl with pbl while at it.

    Fixes: 303ae1cdfdf7 ("rdma/siw: application interface")
    Signed-off-by: Guoqing Jiang <guoqing.jiang@linux.dev>
    Link: https://lore.kernel.org/r/20230821133255.31111-3-guoqing.jiang@linux.dev
    Acked-by: Bernard Metzler <bmt@zurich.ibm.com>
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f7517711ae39cccf994ffc76085f429b4bebdfd5
Author: Guoqing Jiang <guoqing.jiang@linux.dev>
Date:   Mon Aug 21 21:32:53 2023 +0800

    RDMA/siw: Balance the reference of cep->kref in the error path

    [ Upstream commit b056327bee09e6b86683d3f709a438ccd6031d72 ]

    The siw_connect can go to err in below after cep is allocated successfully:

    1. If siw_cm_alloc_work returns failure. In this case socket is not
    assoicated with cep so siw_cep_put can't be called by siw_socket_disassoc.
```

```
    We need to call siw_cep_put twice since cep->kref is increased once after
    it was initialized.

    2. If siw_cm_queue_work can't find a work, which means siw_cep_get is not
    called in siw_cm_queue_work, so cep->kref is increased twice by siw_cep_get
    and when associate socket with cep after it was initialized. So we need to
    call siw_cep_put three times (one in siw_socket_disassoc).

    3. siw_send_mpareqrep returns error, this scenario is similar as 2.

    So we need to remove one siw_cep_put in the error path.

    Fixes: 6c52fdc244b5 ("rdma/siw: connection management")
    Signed-off-by: Guoqing Jiang <guoqing.jiang@linux.dev>
    Link: https://lore.kernel.org/r/20230821133255.31111-2-guoqing.jiang@linux.dev
    Acked-by: Bernard Metzler <bmt@zurich.ibm.com>
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3189341e2f609d48f730b18c8bbbf6783233477
Author: Leon Romanovsky <leon@kernel.org>
Date:   Mon Aug 21 10:57:14 2023 +0300

    Revert "IB/isert: Fix incorrect release of isert connection"

    [ Upstream commit dfe261107c080709459c32695847eec96238852b ]

    Commit: 699826f4e30a ("IB/isert: Fix incorrect release of isert connection") is
    causing problems on OPA when DEVICE_REMOVAL is happening.

     ------------[ cut here ]------------
     WARNING: CPU: 52 PID: 2117247 at drivers/infiniband/core/cq.c:359
    ib_cq_pool_cleanup+0xac/0xb0 [ib_core]
     Modules linked in: nfsd nfs_acl target_core_user uio tcm_fc libfc
    scsi_transport_fc tcm_loop target_core_pscsi target_core_iblock target_core_file
    rpcsec_gss_krb5 auth_rpcgss nfsv4 dns_resolver nfs lockd grace fscache netfs
    rfkill rpcrdma rdma_ucm ib_srpt sunrpc ib_isert iscsi_target_mod target_core_mod
    opa_vnic ib_iser libiscsi ib_umad scsi_transport_iscsi rdma_cm ib_ipoib iw_cm
    ib_cm hfi1(-) rdmavt ib_uverbs intel_rapl_msr intel_rapl_common sb_edac ib_core
    x86_pkg_temp_thermal intel_powerclamp coretemp i2c_i801 mxm_wmi rapl iTCO_wdt
    ipmi_si iTCO_vendor_support mei_me ipmi_devintf mei intel_cstate ioatdma
    intel_uncore i2c_smbus joydev pcspkr lpc_ich ipmi_msghandler acpi_power_meter
    acpi_pad xfs libcrc32c sr_mod sd_mod cdrom t10_pi sg crct10dif_pclmul
    crc32_pclmul crc32c_intel drm_kms_helper drm_shmem_helper ahci libahci
    ghash_clmulni_intel igb drm libata dca i2c_algo_bit wmi fuse
     CPU: 52 PID: 2117247 Comm: modprobe Not tainted 6.5.0-rc1+ #1
     Hardware name: Intel Corporation S2600CWR/S2600CW, BIOS
    SE5C610.86B.01.01.0014.121820151719 12/18/2015
     RIP: 0010:ib_cq_pool_cleanup+0xac/0xb0 [ib_core]
     Code: ff 48 8b 43 40 48 8d 7b 40 48 83 e8 40 4c 39 e7 75 b3 49 83
    c4 10 4d 39 fc 75 94 5b 5d 41 5c 41 5d 41 5e 41 5f c3 cc cc cc cc <0f> 0b eb a1
    90 90 90 90 90 90 90 90 90 90 90 90 90 90 90 0f 1f
     RSP: 0018:ffffc10bea13fc80 EFLAGS: 00010206
     RAX: 000000000000010c RBX: ffff9bf5c7e66c00 RCX: 000000008020001d
     RDX: 000000008020001e RSI: fffff175221f9900 RDI: ffff9bf5c7e67640
     RBP: ffff9bf5c7e67600 R08: ffff9bf5c7e64400 R09: 000000008020001d
     R10: 0000000040000000 R11: 0000000000000000 R12: ffff9bee4b1e8a18
     R13: dead000000000122 R14: dead000000000100 R15: ffff9bee4b1e8a38
     FS:  00007ff1e6d38740(0000) GS:ffff9bfd9fb00000(0000) knlGS:0000000000000000
     CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
     CR2: 00005652044ecc68 CR3: 0000000889b5c005 CR4: 00000000001706e0
     Call Trace:
      <TASK>
      ? __warn+0x80/0x130
      ? ib_cq_pool_cleanup+0xac/0xb0 [ib_core]
```

```
       ? report_bug+0x195/0x1a0
       ? handle_bug+0x3c/0x70
       ? exc_invalid_op+0x14/0x70
       ? asm_exc_invalid_op+0x16/0x20
       ? ib_cq_pool_cleanup+0xac/0xb0 [ib_core]
       disable_device+0x9d/0x160 [ib_core]
       __ib_unregister_device+0x42/0xb0 [ib_core]
       ib_unregister_device+0x22/0x30 [ib_core]
       rvt_unregister_device+0x20/0x90 [rdmavt]
       hfi1_unregister_ib_device+0x16/0xf0 [hfi1]
       remove_one+0x55/0x1a0 [hfi1]
       pci_device_remove+0x36/0xa0
       device_release_driver_internal+0x193/0x200
       driver_detach+0x44/0x90
       bus_remove_driver+0x69/0xf0
       pci_unregister_driver+0x2a/0xb0
       hfi1_mod_cleanup+0xc/0x3c [hfi1]
       __do_sys_delete_module.constprop.0+0x17a/0x2f0
       ? exit_to_user_mode_prepare+0xc4/0xd0
       ? syscall_trace_enter.constprop.0+0x126/0x1a0
       do_syscall_64+0x5c/0x90
       ? syscall_exit_to_user_mode+0x12/0x30
       ? do_syscall_64+0x69/0x90
       ? syscall_exit_work+0x103/0x130
       ? syscall_exit_to_user_mode+0x12/0x30
       ? do_syscall_64+0x69/0x90
       ? exc_page_fault+0x65/0x150
       entry_SYSCALL_64_after_hwframe+0x6e/0xd8
      RIP: 0033:0x7ff1e643f5ab
      Code: 73 01 c3 48 8b 0d 75 a8 1b 00 f7 d8 64 89 01 48 83 c8 ff c3
     66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 1e fa b8 b0 00 00 00 0f 05 <48> 3d 01 f0
     ff ff 73 01 c3 48 8b 0d 45 a8 1b 00 f7 d8 64 89 01 48
      RSP: 002b:00007ffec9103cc8 EFLAGS: 00000206 ORIG_RAX: 00000000000000b0
      RAX: ffffffffffffffda RBX: 00005615267fdc50 RCX: 00007ff1e643f5ab
      RDX: 0000000000000000 RSI: 0000000000000800 RDI: 00005615267fdcb8
      RBP: 00005615267fdc50 R08: 0000000000000000 R09: 0000000000000000
      R10: 00007ff1e659eac0 R11: 0000000000000206 R12: 00005615267fdcb8
      R13: 0000000000000000 R14: 00005615267fdcb8 R15: 00007ffec9105ff8
       </TASK>
      ---[ end trace 0000000000000000 ]---

     And...

      restrack: ------------[ cut here ]------------
      infiniband hfi1_0: BUG: RESTRACK detected leak of resources
      restrack: Kernel PD object allocated by ib_isert is not freed
      restrack: Kernel CQ object allocated by ib_core is not freed
      restrack: Kernel QP object allocated by rdma_cm is not freed
      restrack: ------------[ cut here ]------------

     Fixes: 699826f4e30a ("IB/isert: Fix incorrect release of isert connection")
     Reported-by: Dennis Dalessandro <dennis.dalessandro@cornelisnetworks.com>
     Closes: https://lore.kernel.org/all/921cd1d9-2879-f455-1f50-
  0053fe6a6655@cornelisnetworks.com
     Link:
  https://lore.kernel.org/r/a27982d3235005c58f6d321f3fad5eb6e1beaf9e.1692604607.git.leonro@nvi
  dia.com
     Tested-by: Dennis Dalessandro <dennis.dalessandro@cornelisnetworks.com>
     Signed-off-by: Leon Romanovsky <leonro@nvidia.com>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 8b60a706166de5de82314494704c2419e7657bf8
  Author: Peng Fan <peng.fan@nxp.com>
  Date:   Mon Aug 21 10:39:27 2023 +0800
```

```
    amba: bus: fix refcount leak

    [ Upstream commit e312cbdc11305568554a9e18a2ea5c2492c183f3 ]

    commit 5de1540b7bc4 ("drivers/amba: create devices from device tree")
    increases the refcount of of_node, but not releases it in
    amba_device_release, so there is refcount leak. By using of_node_put
    to avoid refcount leak.

    Fixes: 5de1540b7bc4 ("drivers/amba: create devices from device tree")
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20230821023928.3324283-1-peng.fan@oss.nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66acf347804c5127897811b4afbf6a1fee3dd676
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Mon Aug 14 08:02:11 2023 +0200

    Documentation: devices.txt: Fix minors for ttyCPM*

    [ Upstream commit 4b91dcc2f601cc2098b5fead71344704ddcff8b7 ]

    ttyCPM* devices belong to CPM_UART driver at the first place
    and that driver provides 6 ports.

    Fixes: e29c3f81eb89 ("Documentation: devices.txt: reconcile serial/ucc_uart minor
numers")
    Cc: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Reviewed-by: Randy Dunlap <rdunlap@infradead.org>
    Link:
https://lore.kernel.org/r/27d7124cf86157e2a27c2b039e769041994d3f22.1691992627.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 290e5a9166159970864ab927ce7412ad6a18cbc1
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Mon Aug 14 08:02:10 2023 +0200

    Documentation: devices.txt: Remove ttySIOC*

    [ Upstream commit 27681960f05515555441d7bf58d565cbc68137f3 ]

    IOC3 serial driver was removed, remove associated devices
    from documentation.

    Fixes: 9c860e4cf708 ("tty/serial: remove the ioc3_serial driver")
    Cc: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Link:
https://lore.kernel.org/r/f13b5c64f8cb6d8f2357d7be14397676b27ac2a2.1691992627.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf39757bda6719cb0cf565dad4ef4ee03daa217d
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Mon Aug 14 08:02:09 2023 +0200

    Documentation: devices.txt: Remove ttyIOC*

    [ Upstream commit e327fdc262345ca37b358a51ff0c0046ab1b8d15 ]
```

```
    IOC4 serial driver was removed, remove associated devices
    from documentation.

    Fixes: a017ef17cfd8 ("tty/serial: remove the ioc4_serial driver")
    Cc: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Link:
https://lore.kernel.org/r/b5deb1222eb92017f0efe5b5cae127ac11983b3d.1691992627.git.christophe
.leroy@csgroup.eu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37024a5466e557c39b065542357e69a46e73802a
Author: Yi Yang <yiyang13@huawei.com>
Date:   Thu Aug 17 18:54:06 2023 +0800

    serial: tegra: handle clk prepare error in tegra_uart_hw_init()

    [ Upstream commit 5abd01145d0cc6cd1b7c2fe6ee0b9ea0fa13671e ]

    In tegra_uart_hw_init(), the return value of clk_prepare_enable() should
    be checked since it might fail.

    Fixes: e9ea096dd225 ("serial: tegra: add serial driver")
    Signed-off-by: Yi Yang <yiyang13@huawei.com>
    Link: https://lore.kernel.org/r/20230817105406.228674-1-yiyang13@huawei.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 597a62afe7bd217208a88c1c741e0f078a9aff6b
Author: Sam Protsenko <semen.protsenko@linaro.org>
Date:   Wed Aug 16 15:11:23 2023 -0500

    dt-bindings: usb: samsung,exynos-dwc3: Fix Exynos5433 compatible

    [ Upstream commit 26f4f8358d89c0d9972a30abdb3f3a425ef49e38 ]

    The correct compatible for Exynos5433 is "samsung,exynos5433-dwusb3".
    Fix the typo in its usage.

    Signed-off-by: Sam Protsenko <semen.protsenko@linaro.org>
    Fixes: 949ea75b7ba4 ("dt-bindings: usb: samsung,exynos-dwc3: convert to dtschema")
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230816201123.3530-1-semen.protsenko@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab2755534c5c75019488bdaba7d0c5a178a1d388
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Fri Aug 18 12:29:11 2023 +0200

    dt-bindings: usb: samsung,exynos-dwc3: fix order of clocks on Exynos5433

    [ Upstream commit 8d4ff1351801bd646c9fed7aedb9705250f2c87b ]

    The Exynos5433 DTSI had always different order of DWC USB3 controller
    clocks than the binding.  The order in the binding was introduced in the
    commit 949ea75b7ba4 ("dt-bindings: usb: samsung,exynos-dwc3: convert to
    dtschema") converting to DT schema.  The Linux driver does not care
    about order and was always getting clocks by name.  Therefore assume the
    DTS is the preferred order and correct the binding.

    Fixes: 949ea75b7ba4 ("dt-bindings: usb: samsung,exynos-dwc3: convert to dtschema")
    Cc: Sam Protsenko <semen.protsenko@linaro.org>
```

```
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Acked-by: Rob Herring <robh@kernel.org>
    Reviewed-by: Sam Protsenko <semen.protsenko@linaro.org>
    Link: https://lore.kernel.org/r/20230818102911.18388-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 62b1aa9e13bfca6d0cac2055a2e066f87676a827
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Sat Aug 12 01:16:16 2023 +0200

    interconnect: qcom: bcm-voter: Use enable_maks for keepalive voting

    [ Upstream commit 1a70ca71547be051769f0628aa09717694f508f0 ]

    BCMs with an enable_mask expect to only have that specific value written
    to them. The current implementation only works by miracle for BCMs with
    enable mask == BIT(0), as the minimal vote we've been using so far just
    so happens to be equal to that.

    Use the correct value with keepalive voting.

    Fixes: d8630f050d3f ("interconnect: qcom: Add support for mask-based BCMs")
    Reported-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230811-topic-icc_fix_1he-v2-2-0620af8ac133@linaro.org
    Signed-off-by: Georgi Djakov <djakov@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f357929cac6bf88745a994ae04ac754d119b1fa7
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Sat Aug 12 01:16:15 2023 +0200

    interconnect: qcom: bcm-voter: Improve enable_mask handling

    [ Upstream commit a1f4170dec440f023601d57e49227b784074d218 ]

    We don't need all the complex arithmetic for BCMs utilizing enable_mask,
    as all we need to do is to determine whether there's any user (or
    keepalive) asking for it to be on.

    Separate the logic for such BCMs for a small speed boost.

    Suggested-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Reviewed-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230811-topic-icc_fix_1he-v2-1-0620af8ac133@linaro.org
    Signed-off-by: Georgi Djakov <djakov@kernel.org>
    Stable-dep-of: 1a70ca71547b ("interconnect: qcom: bcm-voter: Use enable_maks for
keepalive voting")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 532ba7e4223c7acb5c58bfabb9d7ea3b05958ec4
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Aug 11 19:34:57 2023 +0200

    interconnect: qcom: sm8450: Enable sync_state

    [ Upstream commit 16862f1b2110eca6330e5be6d804e1a08e06a202 ]

    Enable sync_state on sm8450 so that the interconnect votes actually mean
    anything and aren't just pinned to INT_MAX.

    Fixes: fafc114a468e ("interconnect: qcom: Add SM8450 interconnect provider driver")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
```

```
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Link: https://lore.kernel.org/r/20230811-topic-8450_syncstate-v1-1-
69ae5552a18b@linaro.org
    Signed-off-by: Georgi Djakov <djakov@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6c5d7242bcf2154e9576e5eb2a98c6984ca5ea9a
Author: Chengfeng Ye <dg573847474@gmail.com>
Date:   Thu Aug 17 07:47:08 2023 +0000

    scsi: fcoe: Fix potential deadlock on &fip->ctlr_lock

    [ Upstream commit 1a1975551943f681772720f639ff42fbaa746212 ]

    There is a long call chain that &fip->ctlr_lock is acquired by isr
    fnic_isr_msix_wq_copy() under hard IRQ context. Thus other process context
    code acquiring the lock should disable IRQ, otherwise deadlock could happen
    if the IRQ preempts the execution while the lock is held in process context
    on the same CPU.

    [ISR]
    fnic_isr_msix_wq_copy()
     -> fnic_wq_copy_cmpl_handler()
     -> fnic_fcpio_cmpl_handler()
     -> fnic_fcpio_flogi_reg_cmpl_handler()
     -> fnic_flush_tx()
     -> fnic_send_frame()
     -> fcoe_ctlr_els_send()
     -> spin_lock_bh(&fip->ctlr_lock)

    [Process Context]
    1. fcoe_ctlr_timer_work()
     -> fcoe_ctlr_flogi_send()
     -> spin_lock_bh(&fip->ctlr_lock)

    2. fcoe_ctlr_recv_work()
     -> fcoe_ctlr_recv_handler()
     -> fcoe_ctlr_recv_els()
     -> fcoe_ctlr_announce()
     -> spin_lock_bh(&fip->ctlr_lock)

    3. fcoe_ctlr_recv_work()
     -> fcoe_ctlr_recv_handler()
     -> fcoe_ctlr_recv_els()
     -> fcoe_ctlr_flogi_retry()
     -> spin_lock_bh(&fip->ctlr_lock)

    4. -> fcoe_xmit()
     -> fcoe_ctlr_els_send()
     -> spin_lock_bh(&fip->ctlr_lock)

    spin_lock_bh() is not enough since fnic_isr_msix_wq_copy() is a
    hardirq.

    These flaws were found by an experimental static analysis tool I am
    developing for irq-related deadlock.

    The patch fix the potential deadlocks by spin_lock_irqsave() to disable
    hard irq.

    Fixes: 794d98e77f59 ("[SCSI] libfcoe: retry rejected FLOGI to another FCF if possible")
    Signed-off-by: Chengfeng Ye <dg573847474@gmail.com>
    Link: https://lore.kernel.org/r/20230817074708.7509-1-dg573847474@gmail.com
    Reviewed-by: Davidlohr Bueso <dave@stgolabs.net>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
```

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

commit 124a415f9ba17619b431c75998bfde7b5e2061ce
Author: Tony Battersby <tonyb@cybernetics.com>
Date:   Mon Aug 14 10:03:25 2023 -0400

    scsi: core: Use 32-bit hostnum in scsi_host_lookup()

    [ Upstream commit 62ec2092095b678ff89ce4ba51c2938cd1e8e630 ]

    Change scsi_host_lookup() hostnum argument type from unsigned short to
    unsigned int to match the type used everywhere else.

    Fixes: 6d49f63b415c ("[SCSI] Make host_no an unsigned int")
    Signed-off-by: Tony Battersby <tonyb@cybernetics.com>
    Link: https://lore.kernel.org/r/a02497e7-c12b-ef15-47fc-3f0a0b00ffce@cybernetics.com
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 782c5702b933477b088e80e6d07b9493145b2916
Author: Christopher Bednarz <christopher.n.bednarz@intel.com>
Date:   Fri Aug 18 09:48:38 2023 -0500

    RDMA/irdma: Prevent zero-length STAG registration

    [ Upstream commit bb6d73d9add68ad270888db327514384dfa44958 ]

    Currently irdma allows zero-length STAGs to be programmed in HW during
    the kernel mode fast register flow. Zero-length MR or STAG registration
    disable HW memory length checks.

    Improve gaps in bounds checking in irdma by preventing zero-length STAG or
    MR registrations except if the IB_PD_UNSAFE_GLOBAL_RKEY is set.

    This addresses the disclosure CVE-2023-25775.

    Fixes: b48c24c2d710 ("RDMA/irdma: Implement device supported verb APIs")
    Signed-off-by: Christopher Bednarz <christopher.n.bednarz@intel.com>
    Signed-off-by: Shiraz Saleem <shiraz.saleem@intel.com>
    Link: https://lore.kernel.org/r/20230818144838.1758-1-shiraz.saleem@intel.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4123f4aa50cd9713934b015e451a6e84d9c642c3
Author: Chao Yu <chao@kernel.org>
Date:   Sun Jul 30 22:25:52 2023 +0800

    Revert "f2fs: do not issue small discard commands during checkpoint"

    [ Upstream commit 005abf9e5e0d4dcfce318ae5dbcac32b7bf6b647 ]

    Previously, we have two mechanisms to cache & submit small discards:

    a) set max small discard number in /sys/fs/f2fs/vdb/max_small_discards,
    and checkpoint will cache small discard candidates w/ configured maximum
    number.

    b) call FITRIM ioctl, also, checkpoint in f2fs_trim_fs() will cache small
    discard candidates w/ configured discard granularity, but w/o limitation
    of number. FSTRIM interface is asynchronized, so it won't submit discard
    directly.

    Finally, discard thread will submit them in background periodically.
```

```
    However, after commit 9ac00e7cef10 ("f2fs: do not issue small discard
    commands during checkpoint"), the mechanism a) is broken, since no matter
    how we configure the sysfs entry /sys/fs/f2fs/vdb/max_small_discards,
    checkpoint will not cache small discard candidates any more.

    echo 0 > /sys/fs/f2fs/vdb/max_small_discards
    xfs_io -f /mnt/f2fs/file -c "pwrite 0 2m" -c "fsync"
    xfs_io /mnt/f2fs/file -c "fpunch 0 4k"
    sync
    cat /proc/fs/f2fs/vdb/discard_plist_info |head -2

    echo 100 > /sys/fs/f2fs/vdb/max_small_discards
    rm /mnt/f2fs/file
    xfs_io -f /mnt/f2fs/file -c "pwrite 0 2m" -c "fsync"
    xfs_io /mnt/f2fs/file -c "fpunch 0 4k"
    sync
    cat /proc/fs/f2fs/vdb/discard_plist_info |head -2

    Before the patch:
    Discard pend list(Show diacrd_cmd count on each entry, .:not exist):
      0         .       .       .       .       .       .       .
    Discard pend list(Show diacrd_cmd count on each entry, .:not exist):
      0         3       1       .       .       .       .       .       .

    After the patch:
    Discard pend list(Show diacrd_cmd count on each entry, .:not exist):
      0         .       .       .       .       .       .       .
    Discard pend list(Show diacrd_cmd count on each entry, .:not exist):
      0         .       .       .       .       .       .       .       .

    This patch reverts commit 9ac00e7cef10 ("f2fs: do not issue small discard
    commands during checkpoint") in order to fix this issue.

    Fixes: 9ac00e7cef10 ("f2fs: do not issue small discard commands during checkpoint")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fa508fe42c8b9422026b5d3a1b3a01e12afebe53
Author: Junhao He <hejunhao3@huawei.com>
Date:   Fri Aug 18 16:40:52 2023 +0800

    coresight: trbe: Fix TRBE potential sleep in atomic context

    [ Upstream commit c0a232f1e19e378c5c4e5973a996392942c80090 ]

    smp_call_function_single() will allocate an IPI interrupt vector to
    the target processor and send a function call request to the interrupt
    vector. After the target processor receives the IPI interrupt, it will
    execute arm_trbe_remove_coresight_cpu() call request in the interrupt
    handler.

    According to the device_unregister() stack information, if other process
    is useing the device, the down_write() may sleep, and trigger deadlocks
    or unexpected errors.

      arm_trbe_remove_coresight_cpu
        coresight_unregister
          device_unregister
            device_del
              kobject_del
                __kobject_del
                  sysfs_remove_dir
                    kernfs_remove
                      down_write ---------> it may sleep
```

```
    Add a helper arm_trbe_disable_cpu() to disable TRBE precpu irq and reset
    per TRBE.
    Simply call arm_trbe_remove_coresight_cpu() directly without useing the
    smp_call_function_single(), which is the same as registering the TRBE
    coresight device.

    Fixes: 3fbf7f011f24 ("coresight: sink: Add TRBE driver")
    Signed-off-by: Junhao He <hejunhao3@huawei.com>
    Link: https://lore.kernel.org/r/20230814093813.19152-2-hejunhao3@huawei.com
    [ Remove duplicate cpumask checks during removal ]
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    [ v3 - Remove the operation of assigning NULL to cpudata->drvdata ]
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230818084052.10116-1-hejunhao3@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1b60e7c9fee34eaedf1fc4e0471f75b33f83a4a
Author: Junhao He <hejunhao3@huawei.com>
Date:   Thu Aug 17 16:59:36 2023 +0800

    coresight: Fix memory leak in acpi_buffer->pointer

    [ Upstream commit 1a9e02673e2550f5612099e64e8761f0c8fc0f50 ]

    There are memory leaks reported by kmemleak:
    ...
    unreferenced object 0xffff00213c141000 (size 1024):
      comm "systemd-udevd", pid 2123, jiffies 4294909467 (age 6062.160s)
      hex dump (first 32 bytes):
        04 00 00 00 02 00 00 00 18 10 14 3c 21 00 ff ff  ...........<!...
        00 00 00 00 00 00 00 00 03 00 00 00 10 00 00 00  ................
      backtrace:
        [<000000004b7c9001>] __kmem_cache_alloc_node+0x2f8/0x348
        [<00000000b0fc7ceb>] __kmalloc+0x58/0x108
        [<0000000064ff4695>] acpi_os_allocate+0x2c/0x68
        [<000000007d57d116>] acpi_ut_initialize_buffer+0x54/0xe0
        [<0000000024583908>] acpi_evaluate_object+0x388/0x438
        [<0000000017b2e72b>] acpi_evaluate_object_typed+0xe8/0x240
        [<000000005df0eac2>] coresight_get_platform_data+0x1b4/0x988 [coresight]
    ...

    The ACPI buffer memory (buf.pointer) should be freed. But the buffer
    is also used after returning from acpi_get_dsd_graph().
    Move the temporary variables buf to acpi_coresight_parse_graph(),
    and free it before the function return to prevent memory leak.

    Fixes: 76ffa5ab5b79 ("coresight: Support for ACPI bindings")
    Signed-off-by: Junhao He <hejunhao3@huawei.com>
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230817085937.55590-2-hejunhao3@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit be6eff2a7eeba9c4cf13d60ddd2ba65852393b9d
Author: Suzuki K Poulose <suzuki.poulose@arm.com>
Date:   Mon Jul 10 11:54:59 2023 +0530

    coresight: platform: acpi: Ignore the absence of graph

    [ Upstream commit 3a2888aa1f962c55ca36119aebe67355c7bf54e4 ]

    Some components may not have graph connections for describing
    the trace path. e.g., ETE, where it could directly use the per
    CPU TRBE. Ignore the absence of graph connections
```

```
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Signed-off-by: Anshuman Khandual <anshuman.khandual@arm.com>
    Link: https://lore.kernel.org/r/20230710062500.45147-6-anshuman.khandual@arm.com
    Stable-dep-of: 1a9e02673e25 ("coresight: Fix memory leak in acpi_buffer->pointer")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 60075c64c0d47f22a59cd9df656f3e1a57c4a2b7
Author: Suzuki K Poulose <suzuki.poulose@arm.com>
Date:   Wed Aug 16 13:51:50 2023 +0100

    coresight: trbe: Allocate platform data per device

    [ Upstream commit 39744738a67de9153d73e11817937c0004feab2e ]

    Coresight TRBE driver shares a single platform data (which is empty btw).
    However, with the commit 4e8fe7e5c3a5
    ("coresight: Store pointers to connections rather than an array of them")
    the coresight core would free up the pdata, resulting in multiple attempts
    to free the same pdata for TRBE instances. Fix this by allocating a pdata per
    coresight_device.

    Fixes: 4e8fe7e5c3a5 ("coresight: Store pointers to connections rather than an array of
them")
    Link: https://lore.kernel.org/r/20230814093813.19152-3-hejunhao3@huawei.com
    Reported-by: Junhao He <hejunhao3@huawei.com>
    Cc: Anshuman Khandual <anshuman.khandual@arm.com>
    Cc: James Clark <james.clark@arm.com>
    Tested-by: Junhao He <hejunhao3@huawei.com>
    Link: https://lore.kernel.org/r/20230816141008.535450-2-suzuki.poulose@arm.com
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fc2545bf002a6cfb5064f51c7c5c5280c44a9068
Author: Lu Jialin <lujialin4@huawei.com>
Date:   Thu Aug 10 11:25:28 2023 +0000

    cgroup:namespace: Remove unused cgroup_namespaces_init()

    [ Upstream commit 82b90b6c5b38e457c7081d50dff11ecbafc1e61a ]

    cgroup_namspace_init() just return 0. Therefore, there is no need to
    call it during start_kernel. Just remove it.

    Fixes: a79a908fd2b0 ("cgroup: introduce cgroup namespaces")
    Signed-off-by: Lu Jialin <lujialin4@huawei.com>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d1df07477a8b927c046a82a59368de1d668fb311
Author: Chao Yu <chao@kernel.org>
Date:   Tue Aug 8 08:59:49 2023 +0800

    f2fs: fix to account cp stats correctly

    [ Upstream commit eb61c2cca2eb2110cc7b61a7bc15b3850977a778 ]

    cp_foreground_calls sysfs entry shows total CP call count rather than
    foreground CP call count, fix it.

    Fixes: fc7100ea2a52 ("f2fs: Add f2fs stats to sysfs")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 49dd8d3a3e03c4b705dbf963dacf5f704fe1ead2
Author: Chao Yu <chao@kernel.org>
Date:   Tue Aug 8 08:59:48 2023 +0800

    f2fs: fix to account gc stats correctly

    [ Upstream commit 9bf1dcbdfdc8892d9cfeaeab02519c0ecf17fe51 ]

    As reported, status debugfs entry shows inconsistent GC stats as below:

    GC calls: 6008 (BG: 6161)
      - data segments : 3053 (BG: 3053)
      - node segments : 2955 (BG: 2955)

    Total GC calls is larger than BGGC calls, the reason is:
    - f2fs_stat_info.call_count accounts total migrated section count
    by f2fs_gc()
    - f2fs_stat_info.bg_gc accounts total call times of f2fs_gc() from
    background gc_thread

    Another issue is gc_foreground_calls sysfs entry shows total GC call
    count rather than FGGC call count.

    This patch changes as below for fix:
    - account GC calls and migrated segment count separately
    - support to account migrated section count if it enables large section
    mode
    - fix to show correct value in gc_foreground_calls sysfs entry

    Fixes: fc7100ea2a52 ("f2fs: Add f2fs stats to sysfs")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bbb3cd66301ef752fae2922452660f228d69bcaf
Author: Chao Yu <chao@kernel.org>
Date:   Thu Jul 20 19:29:53 2023 +0800

    Revert "f2fs: fix to do sanity check on extent cache correctly"

    [ Upstream commit 958ccbbf1ce716d77c7cfa79ace50a421c1eed73 ]

    syzbot reports a f2fs bug as below:

    UBSAN: array-index-out-of-bounds in fs/f2fs/f2fs.h:3275:19
    index 1409 is out of range for type '__le32[923]' (aka 'unsigned int[923]')
    Call Trace:
     __dump_stack lib/dump_stack.c:88 [inline]
     dump_stack_lvl+0x1e7/0x2d0 lib/dump_stack.c:106
     ubsan_epilogue lib/ubsan.c:217 [inline]
     __ubsan_handle_out_of_bounds+0x11c/0x150 lib/ubsan.c:348
     inline_data_addr fs/f2fs/f2fs.h:3275 [inline]
     __recover_inline_status fs/f2fs/inode.c:113 [inline]
     do_read_inode fs/f2fs/inode.c:480 [inline]
     f2fs_iget+0x4730/0x48b0 fs/f2fs/inode.c:604
     f2fs_fill_super+0x640e/0x80c0 fs/f2fs/super.c:4601
     mount_bdev+0x276/0x3b0 fs/super.c:1391
     legacy_get_tree+0xef/0x190 fs/fs_context.c:611
     vfs_get_tree+0x8c/0x270 fs/super.c:1519
     do_new_mount+0x28f/0xae0 fs/namespace.c:3335
     do_mount fs/namespace.c:3675 [inline]
     __do_sys_mount fs/namespace.c:3884 [inline]
     __se_sys_mount+0x2d9/0x3c0 fs/namespace.c:3861
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80
```

```
      entry_SYSCALL_64_after_hwframe+0x63/0xcd

    The issue was bisected to:

    commit d48a7b3a72f121655d95b5157c32c7d555e44c05
    Author: Chao Yu <chao@kernel.org>
    Date:   Mon Jan 9 03:49:20 2023 +0000

        f2fs: fix to do sanity check on extent cache correctly

    The root cause is we applied both v1 and v2 of the patch, v2 is the right
    fix, so it needs to revert v1 in order to fix reported issue.

    v1:
    commit d48a7b3a72f1 ("f2fs: fix to do sanity check on extent cache correctly")
    https://lore.kernel.org/lkml/20230109034920.492914-1-chao@kernel.org/

    v2:
    commit 269d11948100 ("f2fs: fix to do sanity check on extent cache correctly")
    https://lore.kernel.org/lkml/20230207134808.1827869-1-chao@kernel.org/

    Reported-by: syzbot+601018296973a481f302@syzkaller.appspotmail.com
    Closes: https://lore.kernel.org/linux-f2fs-
devel/000000000000fcf0690600e4d04d@google.com/
    Fixes: d48a7b3a72f1 ("f2fs: fix to do sanity check on extent cache correctly")
    Signed-off-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fe3545e53d662a5bc4dbef7a9cdf214f4894fb69
Author: Chunhai Guo <guochunhai@vivo.com>
Date:   Thu Aug 3 22:28:42 2023 +0800

    f2fs: Only lfs mode is allowed with zoned block device feature

    [ Upstream commit 2bd4df8fcbc72f58ce3c62ed021ab291ca42de0b ]

    Now f2fs support four block allocation modes: lfs, adaptive,
    fragment:segment, fragment:block. Only lfs mode is allowed with zoned block
    device feature.

    Fixes: 6691d940b0e0 ("f2fs: introduce fragment allocation mode mount option")
    Signed-off-by: Chunhai Guo <guochunhai@vivo.com>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 76a33d238a41bc275d58490b02d8e9db5673d486
Author: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
Date:   Fri Aug 4 18:15:56 2023 +0900

    f2fs: check zone type before sending async reset zone command

    [ Upstream commit 3cb88bc15937990177df1f7eac6f22ebbed19312 ]

    The commit 25f9080576b9 ("f2fs: add async reset zone command support")
    introduced "async reset zone commands" by calling
    __submit_zone_reset_cmd() in async discard operations. However,
    __submit_zone_reset_cmd() is called regardless of zone type of discard
    target zone. When devices have conventional zones, zone reset commands
    are sent to the conventional zones and cause I/O errors.

    Avoid the I/O errors by checking that the discard target zone type is
    sequential write required. If not, handle the discard operation in same
    manner as non-zoned, regular block devices. For that purpose, add a new
    helper function f2fs_bdev_index() which gets index of the zone reset
```

```
      target device.

      Fixes: 25f9080576b9 ("f2fs: add async reset zone command support")
      Signed-off-by: Shin'ichiro Kawasaki <shinichiro.kawasaki@wdc.com>
      Reviewed-by: Chao Yu <chao@kernel.org>
      Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c1abedb8e6b4011858f72f413ea0d834ef427c89
  Author: Christoph Hellwig <hch@lst.de>
  Date:   Fri Jul 7 10:31:49 2023 +0200

      f2fs: don't reopen the main block device in f2fs_scan_devices

      [ Upstream commit 51bf8d3c81992ae57beeaf22df78ed7c2782af9d ]

      f2fs_scan_devices reopens the main device since the very beginning, which
      has always been useless, and also means that we don't pass the right
      holder for the reopen, which now leads to a warning as the core super.c
      holder ops aren't passed in for the reopen.

      Fixes: 3c62be17d4f5 ("f2fs: support multiple devices")
      Fixes: 0718afd47f70 ("block: introduce holder ops")
      Signed-off-by: Christoph Hellwig <hch@lst.de>
      Reviewed-by: Chao Yu <chao@kernel.org>
      Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit b1b5ed4cb11ffd65d85f630aa6745caed94f240e
  Author: Chao Yu <chao@kernel.org>
  Date:   Thu Jul 6 10:06:14 2023 +0800

      f2fs: fix to avoid mmap vs set_compress_option case

      [ Upstream commit b5ab3276eb69cacf44ecfb11b2bfab73096ff4e4 ]

      Compression option in inode should not be changed after they have
      been used, however, it may happen in below race case:

      Thread A                                Thread B
      - f2fs_ioc_set_compress_option
       - check f2fs_is_mmap_file()
       - check get_dirty_pages()
       - check F2FS_HAS_BLOCKS()
                                              - f2fs_file_mmap
                                               - set_inode_flag(FI_MMAP_FILE)
                                              - fault
                                               - do_page_mkwrite
                                                - f2fs_vm_page_mkwrite
                                                - f2fs_get_block_locked
                                               - fault_dirty_shared_page
                                                - set_page_dirty
       - update i_compress_algorithm
       - update i_log_cluster_size
       - update i_cluster_size

      Avoid such race condition by covering f2fs_file_mmap() w/ i_sem lock,
      meanwhile add mmap file check condition in f2fs_may_compress() as well.

      Fixes: e1e8debec656 ("f2fs: add F2FS_IOC_SET_COMPRESS_OPTION ioctl")
      Signed-off-by: Chao Yu <chao@kernel.org>
      Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 4b7c6720f543559b6057049639f028562e8fc57f
```

```
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Jul 9 22:23:24 2023 -0700

    f2fs: fix spelling in ABI documentation

    [ Upstream commit c709d099a0d2befa2b16c249ef8df722b43e6c28 ]

    Correct spelling problems as identified by codespell.

    Fixes: 9e615dbba41e ("f2fs: add missing description for ipu_policy node")
    Fixes: b2e4a2b300e5 ("f2fs: expose discard related parameters in sysfs")
    Fixes: 846ae671ad36 ("f2fs: expose extension_list sysfs entry")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Jaegeuk Kim <jaegeuk@kernel.org>
    Cc: Chao Yu <chao@kernel.org>
    Cc: linux-f2fs-devel@lists.sourceforge.net
    Cc: Yangtao Li <frank.li@vivo.com>
    Cc: Konstantin Vyshetsky <vkon@google.com>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9df818730c3fd6d2aa331a8ff2fb0be201b0c72b
Author: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
Date:   Thu Aug 10 15:33:37 2023 +0200

    media: i2c: rdacm21: Fix uninitialized value

    [ Upstream commit 33c7ae8f49e3413c81e879e1fdfcea4c5516e37b ]

    Fix the following smatch warning:

    drivers/media/i2c/rdacm21.c:373 ov10640_check_id() error: uninitialized
    symbol 'val'.

    Initialize 'val' to 0 in the ov10640_check_id() function.

    Fixes: 2b821698dc73 ("media: i2c: rdacm21: Power up OV10640 before OV490")
    Reported-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9fd2721ed9fee16f6ac49b85e3c2a2ea8df5b91
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:23 2023 +0200

    media: ov2680: Fix regulators being left enabled on ov2680_power_on() errors

    [ Upstream commit 84b4bd7e0d98166aa32fd470e672721190492eae ]

    When the ov2680_power_on() "sensor soft reset failed" path is hit during
    probe() the WARN() about putting an enabled regulator at
    drivers/regulator/core.c:2398 triggers 3 times (once for each regulator),
    filling dmesg with backtraces.

    Fix this by properly disabling the regulators on ov2680_power_on() errors.

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 04593f2d9e6f32f5e48a51b1fe485d2b5a79f85e
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:22 2023 +0200

    media: ov2680: Fix ov2680_set_fmt() which == V4L2_SUBDEV_FORMAT_TRY not working

    [ Upstream commit c0e97a4b4f20639f74cd5809b42ba6cbf9736a7d ]

    ov2680_set_fmt() which == V4L2_SUBDEV_FORMAT_TRY was getting
    the try_fmt v4l2_mbus_framefmt struct from the passed in sd_state
    and then storing the contents of that into the return by reference
    format->format struct.

    While the right thing to do would be filling format->format based on
    the just looked up mode and then store the results of that in
    sd_state->pads[0].try_fmt .

    Before the previous change introducing ov2680_fill_format() this
    resulted in ov2680_set_fmt() which == V4L2_SUBDEV_FORMAT_TRY always
    returning the zero-ed out sd_state->pads[0].try_fmt in format->format
    breaking callers using this.

    After the introduction of ov2680_fill_format() which at least
    initializes sd_state->pads[0].try_fmt properly, format->format
    is now always being filled with the default 800x600 mode set by
    ov2680_init_cfg() independent of the actual requested mode.

    Move the filling of format->format with ov2680_fill_format() to
    before the if (which == V4L2_SUBDEV_FORMAT_TRY) and then store
    the filled in format->format in sd_state->pads[0].try_fmt to
    fix this.

    Note this removes the fmt local variable because IMHO having a local
    variable which points to a sub-struct of one of the function arguments
    just leads to confusion when reading the code.

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c1735ac30c704de32fb94b71b02cf7897e7e12f2
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:21 2023 +0200

    media: ov2680: Add ov2680_fill_format() helper function

    [ Upstream commit 6d6849b2203f3244b575ba01d3e41ee19aa2cadf ]

    Add a ov2680_fill_format() helper function and use this everywhere were
    a v4l2_mbus_framefmt struct needs to be filled in so that the driver always
    fills it consistently.

    This is a preparation patch for fixing ov2680_set_fmt()
    which == V4L2_SUBDEV_FORMAT_TRY calls not properly filling in
    the passed in v4l2_mbus_framefmt struct.

    Note that for ov2680_init_cfg() this now simply always fills
    the try_fmt struct of the passed in sd_state. This is correct because
    ov2680_init_cfg() is never called with a NULL sd_state so the old
    sd_state check is not necessary.
```

```
    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ce53b19673630f632902ea508eca2b4b555135e
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:20 2023 +0200

    media: ov2680: Don't take the lock for try_fmt calls

    [ Upstream commit e521b9cc1a49de677f4fc65909ce4877fbf7b113 ]

    On ov2680_set_fmt() calls with format->which == V4L2_SUBDEV_FORMAT_TRY,
    ov2680_set_fmt() does not talk to the sensor.

    So in this case there is no need to lock the sensor->lock mutex or
    to check that the sensor is streaming.

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 21e63014fb39b8c4935bd4771c95d0f7b18212b9
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:19 2023 +0200

    media: ov2680: Remove VIDEO_V4L2_SUBDEV_API ifdef-s

    [ Upstream commit 49c282d5a8c5f4d1d9088622bec792294c716010 ]

    VIDEO_V4L2_SUBDEV_API is now automatically selected in Kconfig
    for all sensor drivers. Remove the ifdef CONFIG_VIDEO_V4L2_SUBDEV_API
    checks.

    This is a preparation patch for fixing ov2680_set_fmt()
    which == V4L2_SUBDEV_FORMAT_TRY calls not properly filling in
    the passed in v4l2_mbus_framefmt struct.

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 75350ac13ad9d1e8e059ca7f8172d87858c81463
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:18 2023 +0200

    media: ov2680: Fix vflip / hflip set functions

    [ Upstream commit d5d08ad330c9ccebc5e066fda815423a290f48b0 ]

    ov2680_vflip_disable() / ov2680_hflip_disable() pass BIT(0) instead of
    0 as value to ov2680_mod_reg().
```

```
    While fixing this also:

    1. Stop having separate enable/disable functions for hflip / vflip
    2. Move the is_streaming check, which is unique to hflip / vflip
       into the ov2680_set_?flip() functions.

    for a nice code cleanup.

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c273bc365e71f991b3dd2411d186d900f0a9bf6f
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:17 2023 +0200

    media: ov2680: Fix ov2680_bayer_order()

    [ Upstream commit 50a7bad4e0a37d7018ab6fe843dd84bc6b2ecf72 ]

    The index into ov2680_hv_flip_bayer_order[] should be 0-3, but
    ov2680_bayer_order() was using 0 + BIT(2) + (BIT(2) << 1) as
    max index, while the intention was to use: 0 + 1 + 2 as max index.

    Fix the index calculation in ov2680_bayer_order(), while at it
    also just use the ctrl values rather then reading them back using
    a slow i2c-read transaction.

    This also allows making the function void, since there now are
    no more i2c-reads to error check.

    Note the check for the ctrls being NULL is there to allow
    adding an ov2680_fill_format() helper later, which will call
    ov2680_set_bayer_order() during probe() before the ctrls are created.

    [Sakari Ailus: Change all users of ov2680_set_bayer_order() here]

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41db861702265fa0f25efbb99f6bb444c8b9a2c9
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Thu Aug 3 11:33:16 2023 +0200

    media: ov2680: Remove auto-gain and auto-exposure controls

    [ Upstream commit 7b5a42e6ae71927359ea67a2c22570ba97fa4059 ]

    Quoting the OV2680 datasheet:

    "3.2 exposure and gain control

    In the OV2680, the exposure time and gain are set manually from an external
    controller. The OV2680 supports manual gain and exposure control only for
    normal applications, no auto mode."
```

```
    And indeed testing with the atomisp_ov2680 fork of ov2680.c has shown that
    auto-exposure and auto-gain do not work.

    Note that the code setting the auto-exposure flag was broken, callers
    of ov2680_exposure_set() were directly passing !!ctrls->auto_exp->val as
    "bool auto_exp" value, but ctrls->auto_exp is a menu control with:

    enum  v4l2_exposure_auto_type {
            V4L2_EXPOSURE_AUTO = 0,
            V4L2_EXPOSURE_MANUAL = 1,
            ...

    So instead of passing !!ctrls->auto_exp->val they should have been passing
    ctrls->auto_exp->val == V4L2_EXPOSURE_AUTO, iow the passed value was
    inverted of what it should have been.

    Also remove ov2680_g_volatile_ctrl() since without auto support the gain
    and exposure controls are not volatile.

    This also fixes the control values not being properly applied in
    ov2680_mode_set(). The 800x600 mode register-list also sets gain,
    exposure and vflip overriding the last set ctrl values.

    ov2680_mode_set() does call ov2680_gain_set() and ov2680_exposure_set()
    but did this before writing the mode register-list, so these values
    would still be overridden by the mode register-list.

    Add a v4l2_ctrl_handler_setup() call after writing the mode register-list
    to restore all ctrl values. Also remove the ctrls->gain->is_new check from
    ov2680_gain_set() so that the gain always gets restored properly.

    Last since ov2680_mode_set() now calls v4l2_ctrl_handler_setup(), remove
    the v4l2_ctrl_handler_setup() call after ov2680_mode_restore() since
    ov2680_mode_restore() calls ov2680_mode_set().

    Fixes: 3ee47cad3e69 ("media: ov2680: Add Omnivision OV2680 sensor driver")
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Acked-by: Rui Miguel Silva <rmfrfs@gmail.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d77f484453a8fc412ac42cb3e027a439d2e1b4fb
Author: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
Date:   Thu Jul 20 10:08:28 2023 +0200

    media: Documentation: Fix [GS]_ROUTING documentation

    [ Upstream commit 997a6b01cd97b74684728d5af6511c333f25957d ]

    Add mention that successful VIDIOC_SUBDEV_G_ROUTING call will update
    'num_routes' and remove mention about non-existing streams, which is
    incorrect.

    Fixes: ea73eda50813 ("media: Documentation: Add GS_ROUTING documentation")
    Signed-off-by: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 43126fbe50847909dc5a89a696da8caa5e58cf66
Author: Marek Vasut <marex@denx.de>
Date:   Tue Jul 25 00:21:16 2023 +0200
```

```
    media: ov5640: Fix initial RESETB state and annotate timings

    [ Upstream commit a210df337c5f5c2cd82f36c9dbb154ec63365c80 ]

    The initial state of RESETB input signal of OV5640 should be
    asserted, i.e. the sensor should be in reset. This is not the
    case, make it so.

    Since the subsequent assertion of RESETB signal is no longer
    necessary and the timing of the power sequencing could be
    slightly adjusted, add annotations to the delays which match
    OV5640 datasheet rev. 2.03, both:
      figure 2-3 power up timing with internal DVDD
      figure 2-4 power up timing with external DVDD source

    The 5..10ms delay between PWDN assertion and RESETB assertion
    is not even documented in the power sequencing diagram, and
    with this reset fix, it is no longer even necessary.

    Fixes: 19a81c1426c1 ("[media] add Omnivision OV5640 sensor driver")
    Reported-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Signed-off-by: Marek Vasut <marex@denx.de>
    Reviewed-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Tested-by: Jai Luthra <j-luthra@ti.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a595e1d9b321f432ceaff0003347bdb6378dcda
Author: Marek Vasut <marex@denx.de>
Date:   Wed Aug 2 16:47:25 2023 +0200

    media: ov5640: Enable MIPI interface in ov5640_set_power_mipi()

    [ Upstream commit 98cb72d3b9c5e03b10fa993752ecfcbd9c572d8c ]

    Set OV5640_REG_IO_MIPI_CTRL00 bit 2 to 1 instead of 0, since 1 means
    MIPI CSI2 interface, while 0 means CPI parallel interface.

    In the ov5640_set_power_mipi() the interface should obviously be set
    to MIPI CSI2 since this functions is used to power up the sensor when
    operated in MIPI CSI2 mode. The sensor should not be in CPI mode in
    that case.

    This fixes a corner case where capturing the first frame on i.MX8MN
    with CSI/ISI resulted in corrupted frame.

    Fixes: aa4bb8b8838f ("media: ov5640: Re-work MIPI startup sequence")
    Reviewed-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Tested-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com> # [Test on imx6q]
    Signed-off-by: Marek Vasut <marex@denx.de>
    Tested-by: Jai Luthra <j-luthra@ti.com> # [Test on bplay, sk-am62]
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f6d848f21e7983077c32c5cd69bde6f4b225029b
Author: Rahul Rameshbabu <rrameshbabu@nvidia.com>
Date:   Mon Aug 7 09:36:18 2023 -0700

    HID: nvidia-shield: Remove led_classdev_unregister in thunderstrike_create

    [ Upstream commit cb818a047f2b95f3d9e08568ff7f8f513832ff2f ]
```

```
    Avoid calling thunderstrike_led_set_brightness from thunderstrike_create
    when led_classdev_unregister is called. led_classdev_unregister was called
    from thunderstrike_create in the error path. Calling
    thunderstrike_led_set_brightness in this situation is unsafe.

    Fixes: f88af60e74a5 ("HID: nvidia-shield: Support LED functionality for Thunderstrike")
    Signed-off-by: Rahul Rameshbabu <rrameshbabu@nvidia.com>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6360b396e81b5295aa9ee3d4f9af13b9bbbbac65
Author: Illia Ostapyshyn <ostapyshyn@sra.uni-hannover.de>
Date:   Tue Jun 13 17:26:00 2023 +0200

    HID: input: Support devices sending Eraser without Invert

    [ Upstream commit 276e14e6c3993317257e1787e93b7166fbc30905 ]

    Some digitizers (notably XP-Pen Artist 24) do not report the Invert
    usage when erasing.  This causes the device to be permanently stuck with
    the BTN_TOOL_RUBBER tool after sending Eraser, as Invert is the only
    usage that can release the tool.  In this state, Touch and Inrange are
    no longer reported to userspace, rendering the pen unusable.

    Prior to commit 87562fcd1342 ("HID: input: remove the need for
    HID_QUIRK_INVERT"), BTN_TOOL_RUBBER was never set and Eraser events were
    simply translated into BTN_TOUCH without causing an inconsistent state.

    Introduce HID_QUIRK_NOINVERT for such digitizers and detect them during
    hidinput_configure_usage().  This quirk causes the tool to be released
    as soon as Eraser is reported as not set.  Set BTN_TOOL_RUBBER in
    input->keybit when mapping Eraser.

    Fixes: 87562fcd1342 ("HID: input: remove the need for HID_QUIRK_INVERT")
    Co-developed-by: Nils Fuhler <nils@nilsfuhler.de>
    Signed-off-by: Nils Fuhler <nils@nilsfuhler.de>
    Signed-off-by: Illia Ostapyshyn <ostapyshyn@sra.uni-hannover.de>
    Signed-off-by: Jiri Kosina <jkosina@suse.cz>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c8c426fae26086a0ca8ab6cc6da2de79810ec038
Author: David Gow <davidgow@google.com>
Date:   Thu Jul 20 14:45:09 2023 +0200

    drivers: base: Free devm resources when unregistering a device

    [ Upstream commit 699fb50d99039a50e7494de644f96c889279aca3 ]

    In the current code, devres_release_all() only gets called if the device
    has a bus and has been probed.

    This leads to issues when using bus-less or driver-less devices where
    the device might never get freed if a managed resource holds a reference
    to the device. This is happening in the DRM framework for example.

    We should thus call devres_release_all() in the device_del() function to
    make sure that the device-managed actions are properly executed when the
    device is unregistered, even if it has neither a bus nor a driver.

    This is effectively the same change than commit 2f8d16a996da ("devres:
    release resources on device_del()") that got reverted by commit
    a525a3ddeaca ("driver core: free devres in device_release") over
    memory leaks concerns.

    This patch effectively combines the two commits mentioned above to
```

```
    release the resources both on device_del() and device_release() and get
    the best of both worlds.

    Fixes: a525a3ddeaca ("driver core: free devres in device_release")
    Signed-off-by: David Gow <davidgow@google.com>
    Signed-off-by: Maxime Ripard <mripard@kernel.org>
    Link: https://lore.kernel.org/r/20230720-kunit-devm-inconsistencies-test-v3-3-
6aa7e074f373@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5874d0bf813d042e288dc955122ade57cffb3da5
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 11 13:47:04 2023 -0400

    USB: gadget: f_mass_storage: Fix unused variable warning

    [ Upstream commit 55c3e571d2a0aabef4f1354604443f1c415d2e85 ]

    Fix a "variable set but not used" warning in f_mass_storage.c.  rc is
    used if verbose debugging is enabled but not otherwise.

    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Fixes: d5e2b67aae79 ("USB: g_mass_storage: template f_mass_storage.c file created")
    Link: https://lore.kernel.org/r/cfed16c7-aa46-494b-ba84-b0e0dc99be3a@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3ca00640431f3f4f62becf0cad4a2883b8805376
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 11 13:44:38 2023 -0400

    USB: gadget: core: Add missing kerneldoc for vbus_work

    [ Upstream commit 159a98afc88e88f588077afe818081d67f50a5e0 ]

    Add a missing kerneldoc description of the vbus_work field in struct usb_udc.

    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Fixes: 50966da807c8 ("usb: gadget: udc: core: Offload usb_udc_vbus_handler processing")
    Link: https://lore.kernel.org/r/1e5e7cda-b2c8-4917-9952-4354f365ede0@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5e3b4e5838c4ca1fda7da286f7307d9c6d2a34ee
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Jul 9 22:23:05 2023 -0700

    docs: ABI: fix spelling/grammar in SBEFIFO timeout interface

    [ Upstream commit 2cd9ec2a51474d4c0b4d2a061f2de7da34eff476 ]

    Correct spelling problems as identified by codespell.
    Correct one grammar error.

    Fixes: 9a93de620e0a ("docs: ABI: testing: Document the SBEFIFO timeout interface")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Eddie James <eajames@linux.ibm.com>
    Cc: Joel Stanley <joel@jms.id.au>
    Link: https://lore.kernel.org/r/20230710052305.29611-1-rdunlap@infradead.org
    Signed-off-by: Joel Stanley <joel@jms.id.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de031db3708f356a52eca21ebf7e7da944f2dd92
Author: Hans de Goede <hdegoede@redhat.com>
```

```
Date:   Wed Jul 5 23:29:54 2023 +0200

    media: ipu-bridge: Do not use on stack memory for software_node.name field

    [ Upstream commit 11e0a7c8e04ee5f406f2baa27761746cbedcfa11 ]

    Commit 567f97bd381f ("media: ipu3-cio2: support multiple sensors and VCMs
    with same HID") introduced an on stack vcm_name and then uses this for
    the name field of the software_node struct used for the vcm.

    But the software_node struct is much longer lived then the current
    stack-frame, so this is no good.

    Instead extend the ipu_node_names struct with an extra field to store
    the vcm software_node name and use that.

    Note this also changes the length of the allocated buffer from
    ACPI_ID_LEN + 4 to 16. the name is filled with "<ipu_vcm_types[x]>-%u"
    where ipu_vcm_types[x] is not an ACPI_ID. The maximum length of
    the strings in the ipu_vcm_types[] array is 11 + 5 bytes for "-255\0"
    means 16 bytes are needed in the worst case scenario.

    Fixes: 567f97bd381f ("media: ipu3-cio2: support multiple sensors and VCMs with same
HID")
    Cc: Bingbu Cao <bingbu.cao@intel.com>
    Reviewed-by: Andy Shevchenko <andy@kernel.org>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9c96bea0e5f4adabcf8359eb54d67abd71d74be2
Author: Bingbu Cao <bingbu.cao@intel.com>
Date:   Thu May 18 12:05:21 2023 +0200

    media: ipu3-cio2: rename cio2 bridge to ipu bridge and move out of ipu3

    [ Upstream commit 881ca25978c6f536a00205daa8b2452edd057ff9 ]

    cio2 bridge was involved along with IPU3. However, in fact all Intel IPUs
    besides IPU3 CIO2 need this bridge driver. This patch move bridge driver
    out of ipu3 directory and rename as ipu-bridge. Then it can be worked with
    IPU3 and other Intel IPUs.

    Signed-off-by: Bingbu Cao <bingbu.cao@intel.com>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Stable-dep-of: 11e0a7c8e04e ("media: ipu-bridge: Do not use on stack memory for
software_node.name field")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e08b091e33ecf6e4cb2c0c5820a69abe7673280b
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Wed Jul 5 23:29:53 2023 +0200

    media: ipu-bridge: Fix null pointer deref on SSDB/PLD parsing warnings

    [ Upstream commit 284be5693163343e1cf17c03917eecd1d6681bcf ]

    When ipu_bridge_parse_rotation() and ipu_bridge_parse_orientation() run
    sensor->adev is not set yet.

    So if either of the dev_warn() calls about unknown values are hit this
```

```
    will lead to a NULL pointer deref.

    Set sensor->adev earlier, with a borrowed ref to avoid making unrolling
    on errors harder, to fix this.

    Fixes: 485aa3df0dff ("media: ipu3-cio2: Parse sensor orientation and rotation")
    Cc: Fabian Wüthrich <me@fabwu.ch>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Daniel Scally <dan.scally@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e17ff9436ec37b33a905c487ae4ce86a01e38818
Author: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
Date:   Mon Apr 24 12:51:24 2023 +0300

    arm64: defconfig: Drop CONFIG_VIDEO_IMX_MEDIA

    [ Upstream commit 0ca2fbab99b12bb81fceaafe5495c00d76789a37 ]

    CONFIG_VIDEO_IMX_MEDIA isn't needed on arm64 platforms since commit
    9f257f502c2e ("media: imx: Unstage the imx7-media-csi driver") which
    moved the last arm64 driver depending on that Kconfig symbol out of
    staging. Drop it from the arm64 defconfig.

    Fixes: 9f257f502c2e ("media: imx: Unstage the imx7-media-csi driver")
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cb6b069dbc8c013592adc3811989296ba5038f49
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue May 30 14:30:35 2023 +0200

    media: venus: hfi_venus: Only consider sys_idle_indicator on V1

    [ Upstream commit 6283e4834c69fa93a108efa18c6aa09c7e626f49 ]

    As per information from Qualcomm [1], this property is not really
    supported beyond msm8916 (HFI V1) and some newer HFI versions really
    dislike receiving it, going as far as crashing the device.

    Only consider toggling it (via the module option) on HFIV1.
    While at it, get rid of the global static variable (which defaulted
    to zero) which was never explicitly assigned to for V1.

    Note: [1] is a reply to the actual message in question, as lore did not
    properly receive some of the emails..

    [1] https://lore.kernel.org/lkml/955cd520-3881-0c22-d818-13fe9a47e124@linaro.org/
    Fixes: 7ed9e0b3393c ("media: venus: hfi, vdec: v6 Add IS_V6() to existing IS_V4() if
locations")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Stanimir Varbanov <stanimir.k.varbanov@gmail.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 10702a24fc396515da97447720c7cd80390fb596
Author: Colin Ian King <colin.i.king@gmail.com>
Date:   Thu Jul 27 19:40:07 2023 +0200

    media: go7007: Remove redundant if statement

    [ Upstream commit f33cb49081da0ec5af0888f8ecbd566bd326eed1 ]
```

```
    The if statement that compares msgs[i].len != 3 is always false because
    it is in a code block where msg[i].len is equal to 3. The check is
    redundant and can be removed.

    As detected by cppcheck static analysis:
    drivers/media/usb/go7007/go7007-i2c.c:168:20: warning: Opposite inner
    'if' condition leads to a dead code block. [oppositeInnerCondition]

    Link: https://lore.kernel.org/linux-media/20230727174007.635572-1-colin.i.king@gmail.com

    Fixes: 866b8695d67e ("Staging: add the go7007 video driver")
    Signed-off-by: Colin Ian King <colin.i.king@gmail.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0525b6d08b18c53f914bea9543073d1047b1fdf3
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Mon Jun 12 15:58:38 2023 +0200

    media: cec: core: add adap_unconfigured() callback

    [ Upstream commit 948a77aaecf202f722cf2264025f9987e5bd5c26 ]

    The adap_configured() callback was called with the adap->lock mutex
    held if the 'configured' argument was false, and without the adap->lock
    mutex held if that argument was true.

    That was very confusing, and so split this up in a adap_unconfigured()
    callback and a high-level configured() callback.

    This also makes it easier to understand when the mutex is held: all
    low-level adap_* callbacks are called with the mutex held. All other
    callbacks are called without that mutex held.

    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Fixes: f1b57164305d ("media: cec: add optional adap_configured callback")
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a6acb8e7ceb6eaa9c029b2184135ea3f8b65be36
Author: Hans Verkuil <hverkuil-cisco@xs4all.nl>
Date:   Mon Jun 12 15:58:37 2023 +0200

    media: cec: core: add adap_nb_transmit_canceled() callback

    [ Upstream commit da53c36ddd3f118a525a04faa8c47ca471e6c467 ]

    A potential deadlock was found by Zheng Zhang with a local syzkaller
    instance.

    The problem is that when a non-blocking CEC transmit is canceled by calling
    cec_data_cancel, that in turn can call the high-level received() driver
    callback, which can call cec_transmit_msg() to transmit a new message.

    The cec_data_cancel() function is called with the adap->lock mutex held,
    and cec_transmit_msg() tries to take that same lock.

    The root cause is that the received() callback can either be used to pass
    on a received message (and then adap->lock is not held), or to report a
    canceled transmit (and then adap->lock is held).

    This is confusing, so create a new low-level adap_nb_transmit_canceled
    callback that reports back that a non-blocking transmit was canceled.
```

```
    And the received() callback is only called when a message is received,
    as was the case before commit f9d0ecbf56f4 ("media: cec: correctly pass
    on reply results") complicated matters.

    Reported-by: Zheng Zhang <zheng.zhang@email.ucr.edu>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Fixes: f9d0ecbf56f4 ("media: cec: correctly pass on reply results")
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d9e03bec147407826266580e7d6ec427241d859
Author: Armin Wolf <W_Armin@gmx.de>
Date:   Sat Aug 5 07:36:10 2023 +0200

    platform/x86: dell-sysman: Fix reference leak

    [ Upstream commit 7295a996fdab7bf83dc3d4078fa8b139b8e0a1bf ]

    If a duplicate attribute is found using kset_find_obj(),
    a reference to that attribute is returned. This means
    that we need to dispose it accordingly. Use kobject_put()
    to dispose the duplicate attribute in such a case.

    Compile-tested only.

    Fixes: e8a60aa7404b ("platform/x86: Introduce support for Systems Management Driver over
WMI for Dell Systems")
    Signed-off-by: Armin Wolf <W_Armin@gmx.de>
    Link: https://lore.kernel.org/r/20230805053610.7106-1-W_Armin@gmx.de
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1578d9fe75f56596809611547e518deb2f8fffac
Author: Yanfei Xu <yanfei.xu@intel.com>
Date:   Wed Aug 9 20:48:04 2023 +0800

    iommu/vt-d: Fix to flush cache of PASID directory table

    [ Upstream commit 8a3b8e63f8371c1247b7aa24ff9c5312f1a6948b ]

    Even the PCI devices don't support pasid capability, PASID table is
    mandatory for a PCI device in scalable mode. However flushing cache
    of pasid directory table for these devices are not taken after pasid
    table is allocated as the "size" of table is zero. Fix it by
    calculating the size by page order.

    Found this when reading the code, no real problem encountered for now.

    Fixes: 194b3348bdbb ("iommu/vt-d: Fix PASID directory pointer coherency")
    Suggested-by: Lu Baolu <baolu.lu@linux.intel.com>
    Signed-off-by: Yanfei Xu <yanfei.xu@intel.com>
    Link: https://lore.kernel.org/r/20230616081045.721873-1-yanfei.xu@intel.com
    Signed-off-by: Lu Baolu <baolu.lu@linux.intel.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f259e8badc8825e4a5e574a43894d1a7f26a62d8
Author: Vijaya Krishna Nivarthi <quic_vnivarth@quicinc.com>
Date:   Wed Aug 9 16:23:13 2023 +0530

    tty: serial: qcom-geni-serial: Poll primary sequencer irq status after cancel_tx

    [ Upstream commit 9c8441330bb399cba6177acce9b0e68c0dbaa597 ]
```

```
    TX is handled by primary sequencer. After cancelling primary command, poll
    primary sequencer's irq status instead of that of secondary.
    While at it, also remove a couple of redundant lines that read from IRQ_EN
    register and write back same.

    Fixes: 2aaa43c70778 ("tty: serial: qcom-geni-serial: add support for serial engine DMA")
    Signed-off-by: Vijaya Krishna Nivarthi <quic_vnivarth@quicinc.com>
    Link: https://lore.kernel.org/r/1691578393-9891-1-git-send-email-
quic_vnivarth@quicinc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a35d154ae98b7b6969a997f249a95eb7a1def268
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Thu Jun 22 11:27:39 2023 +0200

    iommu/qcom: Disable and reset context bank before programming

    [ Upstream commit 9f3fef23d9b5a858a6e6d5f478bb1b6b76265e76 ]

    Writing the new TTBRs, TCRs and MAIRs on a previously enabled
    context bank may trigger a context fault, resulting in firmware
    driven AP resets: change the domain initialization programming
    sequence to disable the context bank(s) and to also clear the
    related fault address (CB_FAR) and fault status (CB_FSR)
    registers before writing new values to TTBR0/1, TCR/TCR2, MAIR0/1.

    Fixes: 0ae349a0f33f ("iommu/qcom: Add qcom_iommu")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230622092742.74819-4-
angelogioacchino.delregno@collabora.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7786e02c2c605000cb340740a87ea48ac6082c42
Author: Eddie James <eajames@linux.ibm.com>
Date:   Mon Jun 12 14:56:50 2023 -0500

    fsi: aspeed: Reset master errors after CFAM reset

    [ Upstream commit 52300909f4670ac552bfeb33c1355b896eac8c06 ]

    It has been observed that sometimes the FSI master will return all 0xffs
    after a CFAM has been taken out of reset, without presenting any error.
    Resetting the FSI master errors resolves the issue.

    Fixes: 4a851d714ead ("fsi: aspeed: Support CFAM reset GPIO")
    Signed-off-by: Eddie James <eajames@linux.ibm.com>
    Link: https://lore.kernel.org/r/20230612195657.245125-8-eajames@linux.ibm.com
    Signed-off-by: Joel Stanley <joel@jms.id.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2573e4f4da7f52987387b7f0636e9279c556b818
Author: Xiang Yang <xiangyang3@huawei.com>
Date:   Fri Aug 4 10:25:25 2023 +0800

    IB/uverbs: Fix an potential error pointer dereference

    [ Upstream commit 26b7d1a27167e7adf75b150755e05d2bc123ce55 ]

    smatch reports the warning below:
    drivers/infiniband/core/uverbs_std_types_counters.c:110
    ib_uverbs_handler_UVERBS_METHOD_COUNTERS_READ() error: 'uattr'
    dereferencing possible ERR_PTR()
```

```
    The return value of uattr maybe ERR_PTR(-ENOENT), fix this by checking
    the value of uattr before using it.

    Fixes: ebb6796bd397 ("IB/uverbs: Add read counters support")
    Signed-off-by: Xiang Yang <xiangyang3@huawei.com>
    Link: https://lore.kernel.org/r/20230804022525.1916766-1-xiangyang3@huawei.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 412ba891aaa5ce231f0e72fa6a3ee767f088154d
Author: Chengchang Tang <tangchengchang@huawei.com>
Date:   Fri Aug 4 09:27:11 2023 +0800

    RDMA/hns: Fix CQ and QP cache affinity

    [ Upstream commit 9e03dbea2b0634b21a45946b4f8097e0dc86ebe1 ]

    Currently, the affinity between QP cache and CQ cache is not
    considered when assigning QPN, it will affect the message rate of HW.

    Allocate QPN from QP cache with better CQ affinity to get better
    performance.

    Fixes: 71586dd20010 ("RDMA/hns: Create QP with selected QPN for bank load balance")
    Signed-off-by: Chengchang Tang <tangchengchang@huawei.com>
    Signed-off-by: Junxian Huang <huangjunxian6@hisilicon.com>
    Link: https://lore.kernel.org/r/20230804012711.808069-5-huangjunxian6@hisilicon.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a8bcffe56aca3eee4f0ca2c05f8c4595a9d9bea
Author: Junxian Huang <huangjunxian6@hisilicon.com>
Date:   Fri Aug 4 09:27:10 2023 +0800

    RDMA/hns: Fix inaccurate error label name in init instance

    [ Upstream commit c9c0bd3c177d93d80968f720304087ba83fe8f74 ]

    This patch fixes inaccurate error label name in init instance.

    Fixes: 70f92521584f ("RDMA/hns: Use the reserved loopback QPs to free MR before
destroying MPT")
    Signed-off-by: Junxian Huang <huangjunxian6@hisilicon.com>
    Link: https://lore.kernel.org/r/20230804012711.808069-4-huangjunxian6@hisilicon.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 528b1481001860994826047576fa46c3976f572
Author: Junxian Huang <huangjunxian6@hisilicon.com>
Date:   Fri Aug 4 09:27:09 2023 +0800

    RDMA/hns: Fix incorrect post-send with direct wqe of wr-list

    [ Upstream commit 706efac4477cdb8be857f6322457de524acc02ff ]

    Currently, direct wqe is not supported for wr-list. RoCE driver excludes
    direct wqe for wr-list by judging whether the number of wr is 1.

    For a wr-list where the second wr is a length-error atomic wr, the
    post-send driver handles the first wr and adds 1 to the wr number counter
    firstly. While handling the second wr, the driver finds out a length error
    and terminates the wr handle process, remaining the counter at 1. This
    causes the driver mistakenly judges there is only 1 wr and thus enters
    the direct wqe process, carrying the current length-error atomic wqe.
```

```
    This patch fixes the error by adding a judgement whether the current wr
    is a bad wr. If so, use the normal doorbell process but not direct wqe
    despite the wr number is 1.

    Fixes: 01584a5edcc4 ("RDMA/hns: Add support of direct wqe")
    Signed-off-by: Junxian Huang <huangjunxian6@hisilicon.com>
    Link: https://lore.kernel.org/r/20230804012711.808069-3-huangjunxian6@hisilicon.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3c2ec23cb5eb991e0f8ae13a02e628644a58766
Author: Chengchang Tang <tangchengchang@huawei.com>
Date:   Fri Aug 4 09:27:08 2023 +0800

    RDMA/hns: Fix port active speed

    [ Upstream commit df1bcf90a66a10967a3a43510b42cb3566208011 ]

    HW supports a variety of different speed, but the current speed
    is fixed.

    The real speed should be querried from ethernet.

    Fixes: 9a4435375cd1 ("IB/hns: Add driver files for hns RoCE driver")
    Signed-off-by: Chengchang Tang <tangchengchang@huawei.com>
    Signed-off-by: Junxian Huang <huangjunxian6@hisilicon.com>
    Link: https://lore.kernel.org/r/20230804012711.808069-2-huangjunxian6@hisilicon.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20aea02d5a10577f47ade42ef3a880544360687d
Author: Kalesh AP <kalesh-anakkur.purayil@broadcom.com>
Date:   Thu Aug 3 01:45:22 2023 -0700

    RDMA/bnxt_re: Remove a redundant flag

    [ Upstream commit fd28c8a8c7a10e7b53851129c6d8dc5945108fe9 ]

    After the cited commit, BNXT_RE_FLAG_GOT_MSIX is redundant.
    Remove it.

    Fixes: 303432211324 ("bnxt_en: Remove runtime interrupt vector allocation")
    Signed-off-by: Kalesh AP <kalesh-anakkur.purayil@broadcom.com>
    Signed-off-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Link: https://lore.kernel.org/r/1691052326-32143-3-git-send-email-
selvin.xavier@broadcom.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50d09f0a121bdf756547ee21a43c43576ca25310
Author: Kalesh AP <kalesh-anakkur.purayil@broadcom.com>
Date:   Thu Aug 3 01:45:21 2023 -0700

    RDMA/bnxt_re: Fix max_qp count for virtual functions

    [ Upstream commit f19fba1f79dc1fb298de7dcbaae9f6299381aeea ]

    Driver has not accounted QP1 for virtual functions
    when fetching device attributes and hence max_qp
    count is one less than active_qp count. Fixed driver
    so that it counts QP1 for virtual functions as well
    while fetching device attributes

    Fixes: ccd9d0d3dffc ("RDMA/bnxt_re: Enable RoCE on virtual functions")
```

```
    Signed-off-by: Saravanan Vajravel <saravanan.vajravel@broadcom.com>
    Signed-off-by: Kalesh AP <kalesh-anakkur.purayil@broadcom.com>
    Signed-off-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Link: https://lore.kernel.org/r/1691052326-32143-2-git-send-email-
selvin.xavier@broadcom.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 48e7e1bb5219bd4437055bef3116d29ccb9d65cc
Author: Chandramohan Akula <chandramohan.akula@broadcom.com>
Date:   Tue Jul 18 22:02:53 2023 -0700

    RDMA/bnxt_re: Initialize Doorbell pacing feature

    [ Upstream commit 586e613d37ec35572a332839973b9c3bccd0c545 ]

    Checks for pacing feature capability and get the doorbell pacing
    configuration using FW commands. Allocate a page and initialize
    the pacing parameters for the applications. Cleanup the page and
    de-initialize the pacing during device removal.

    Link: https://lore.kernel.org/r/1689742977-9128-4-git-send-email-
selvin.xavier@broadcom.com
    Signed-off-by: Chandramohan Akula <chandramohan.akula@broadcom.com>
    Signed-off-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: f19fba1f79dc ("RDMA/bnxt_re: Fix max_qp count for virtual functions")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d015f4d904814f3dce294540a44c59a02fa17b51
Author: Chandramohan Akula <chandramohan.akula@broadcom.com>
Date:   Tue Jul 18 22:02:52 2023 -0700

    bnxt_en: Share the bar0 address with the RoCE driver

    [ Upstream commit 61220e098e858951f1926d66c1490a96351e1c85 ]

    Add a parameter in the bnxt_en_dev structure to share the bar0 address
    with RoCE driver.

    Link: https://lore.kernel.org/r/1689742977-9128-3-git-send-email-
selvin.xavier@broadcom.com
    CC: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Chandramohan Akula <chandramohan.akula@broadcom.com>
    Signed-off-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: f19fba1f79dc ("RDMA/bnxt_re: Fix max_qp count for virtual functions")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e2479ba775472c5f3fe71175c0508f5b7821d8bc
Author: Chandramohan Akula <chandramohan.akula@broadcom.com>
Date:   Tue Jul 18 22:02:51 2023 -0700

    bnxt_en: Update HW interface headers

    [ Upstream commit cf1694f09894e760f4e2cf068ee6519d11cd0ede ]

    Updating the HW structures for the doorbell pacing related
    information. Newly added interface structures will be used in
    the followup patches.

    Link: https://lore.kernel.org/r/1689742977-9128-2-git-send-email-
selvin.xavier@broadcom.com
    CC: Michael Chan <michael.chan@broadcom.com>
    Signed-off-by: Chandramohan Akula <chandramohan.akula@broadcom.com>
```

```
    Signed-off-by: Selvin Xavier <selvin.xavier@broadcom.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Stable-dep-of: f19fba1f79dc ("RDMA/bnxt_re: Fix max_qp count for virtual functions")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a976c52be3a55df1081c5b90e84cb80d2d80ff6d
Author: Zhu Wang <wangzhu9@huawei.com>
Date:   Mon Jul 31 19:27:58 2023 +0800

    iommu: Remove kernel-doc warnings

    [ Upstream commit 6b7867b5b8a6b14c487bf04a693ab424c7a8718d ]

    Remove kernel-doc warnings:

    drivers/iommu/iommu.c:3261: warning: Function parameter or member 'group'
    not described in 'iommu_group_release_dma_owner'
    drivers/iommu/iommu.c:3261: warning: Excess function parameter 'dev'
    description in 'iommu_group_release_dma_owner'
    drivers/iommu/iommu.c:3275: warning: Function parameter or member 'dev'
    not described in 'iommu_device_release_dma_owner'
    drivers/iommu/iommu.c:3275: warning: Excess function parameter 'group'
    description in 'iommu_device_release_dma_owner'

    Signed-off-by: Zhu Wang <wangzhu9@huawei.com>
    Fixes: 89395ccedbc1 ("iommu: Add device-centric DMA ownership interfaces")
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Link: https://lore.kernel.org/r/20230731112758.214775-1-wangzhu9@huawei.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 937bc6cd9b0d248830cd5003bb824e0a36ce8c95
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date:   Mon Jul 24 14:36:05 2023 -0300

    iommu/sprd: Add missing force_aperture

    [ Upstream commit d48a51286c698f7fe8efc688f23a532f4fe9a904 ]

    force_aperture was intended to false only by GART drivers that have an
    identity translation outside the aperture. This does not describe sprd, so
    add the missing 'force_aperture = true'.

    Fixes: b23e4fc4e3fa ("iommu: add Unisoc IOMMU basic driver")
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Acked-by: Chunyan Zhang <zhang.lyra@gmail.com>
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6fe531f18218d11bdf558b26cd9b285b7cab03c2
Author: Chengci.Xu <chengci.xu@mediatek.com>
Date:   Fri Jun 2 17:02:22 2023 +0800

    iommu/mediatek: Fix two IOMMU share pagetable issue

    [ Upstream commit cf69ef46dbd980a0b1c956d668e066a73e0acd0f ]

    Prepare for mt8188 to fix a two IOMMU HWs share pagetable issue.

    We have two MM IOMMU HWs in mt8188, one is VPP-IOMMU, the other is
    VDO-IOMMU. The 2 MM IOMMU HWs share pagetable don't work in this case:
     a) VPP-IOMMU probe firstly.
     b) VDO-IOMMU probe.
     c) The master for VDO-IOMMU probe (means frstdata is vpp-iommu).
     d) The master in another domain probe. No matter it is vdo or vpp.
```

```
    Then it still create a new pagetable in step d). The problem is
    "frstdata->bank[0]->m4u_dom" was not initialized. Then when d) enter, it
    still create a new one.

    In this patch, we create a new variable "share_dom" for this share
    pgtable case, it should be helpful for readable. and put all the share
    pgtable logic in the mtk_iommu_domain_finalise.

    In mt8195, the master of VPP-IOMMU probes before than VDO-IOMMU
    from its dtsi node sequence, we don't see this issue in it. Prepare for
    mt8188.

    Fixes: 645b87c190c9 ("iommu/mediatek: Fix 2 HW sharing pgtable issue")
    Signed-off-by: Chengci.Xu <chengci.xu@mediatek.com>
    Signed-off-by: Yong Wu <yong.wu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Link: https://lore.kernel.org/r/20230602090227.7264-3-yong.wu@mediatek.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 290f021669d63163149ffb6ad919901c87878983
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Fri Aug 4 15:28:49 2023 +0200

    extcon: cht_wc: add POWER_SUPPLY dependency

    [ Upstream commit d20a3a8a32e3fa564ff25da860c5fc1a97642dfe ]

    The driver fails to link when CONFIG_POWER_SUPPLY is disabled:

    x86_64-linux-ld: vmlinux.o: in function `cht_wc_extcon_psy_get_prop':
    extcon-intel-cht-wc.c:(.text+0x15ccda7): undefined reference to
`power_supply_get_drvdata'
    x86_64-linux-ld: vmlinux.o: in function `cht_wc_extcon_pwrsrc_event':
    extcon-intel-cht-wc.c:(.text+0x15cd3e9): undefined reference to `power_supply_changed'
    x86_64-linux-ld: vmlinux.o: in function `cht_wc_extcon_probe':
    extcon-intel-cht-wc.c:(.text+0x15cd596): undefined reference to
`devm_power_supply_register'

    It should be possible to change the driver to not require this at
    compile time and still provide other functions, but adding a hard
    Kconfig dependency does not seem to have any practical downsides
    and is simpler since the option is normally enabled anyway.

    Fixes: 66e31186cd2aa ("extcon: intel-cht-wc: Add support for registering a power_supply
class-device")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Hans de Goede <hdegoede@redhat.com>
    Signed-off-by: Chanwoo Choi <cw00.choi@samsung.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 568800390f41fdd8733261740cb1f0e80c625d60
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Jul 24 14:18:16 2023 +0200

    kernfs: add stub helper for kernfs_generic_poll()

    [ Upstream commit 79038a99445f69c5d28494dd4f8c6f0509f65b2e ]

    In some randconfig builds, kernfs ends up being disabled, so there is no prototype
    for kernfs_generic_poll()

    In file included from kernel/sched/build_utility.c:97:
    kernel/sched/psi.c:1479:3: error: implicit declaration of function 'kernfs_generic_poll'
```

```
is invalid in C99 [-Werror,-Wimplicit-function-declaration]
                    kernfs_generic_poll(t->of, wait);
                    ^
```

    Add a stub helper for it, as we have it for other kernfs functions.

    Fixes: aff037078ecae ("sched/psi: use kernfs polling functions for PSI trigger polling")
    Fixes: 147e1a97c4a0b ("fs: kernfs: add poll file operation")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Chengming Zhou <zhouchengming@bytedance.com>
    Link: https://lore.kernel.org/r/20230724121823.1357562-1-arnd@kernel.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2117576a33b149cc9fecb29971c6e14c76142e10
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date: Mon Jul 24 14:40:46 2023 -0300

    driver core: Call dma_cleanup() on the test_remove path

    [Upstream commit f429378a9bf84d79a7e2cae05d2e3384cf7d68ba]

    When test_remove is enabled really_probe() does not properly pair
    dma_configure() with dma_remove(), it will end up calling dma_configure()
    twice. This corrupts the owner_cnt and renders the group unusable with
    VFIO/etc.

    Add the missing cleanup before going back to re_probe.

    Fixes: 25f3bcfc54bc ("driver core: Add dma_cleanup callback in bus_type")
    Reported-by: Zenghui Yu <yuzenghui@huawei.com>
    Tested-by: Zenghui Yu <yuzenghui@huawei.com>
    Closes: https://lore.kernel.org/all/6472f254-c3c4-8610-4a37-8d9dfdd54ce8@huawei.com/
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Link: https://lore.kernel.org/r/0-v2-4deed94e283e+40948-
really_probe_dma_cleanup_jgg@nvidia.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 437b2e21a70414def9722c64311de6178e4ced22
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date: Tue Jul 18 10:03:49 2023 +0300

    driver core: test_async: fix an error code

    [Upstream commit 22d2381bbd70a5853c2ee77522f4965139672db9]

    The test_platform_device_register_node() function should return error
    pointers instead of NULL. That is what the callers are expecting.

    Fixes: 57ea974fb871 ("driver core: Rewrite test_async_driver_probe to cover
serialization and NUMA affinity")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Link: https://lore.kernel.org/r/1e11ed19-e1f6-43d8-b352-474134b7c008@moroto.mountain
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9bca9e675a0d8df3246711c4b33728c1fc01b77e
Author: Rob Clark <robdclark@chromium.org>
Date: Mon Jul 24 07:49:41 2023 -0700

    dma-buf/sync_file: Fix docs syntax

    [Upstream commit 05d56d8079d510a2994039470f65bea85f0075ee]

```
    Fixes the warning:

      include/uapi/linux/sync_file.h:77: warning: Function parameter or member 'num_fences'
not described in 'sync_file_info'

    Fixes: 2d75c88fefb2 ("staging/android: refactor SYNC IOCTLs")
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Reviewed-by: Randy Dunlap <rdunlap@infradead.org>
    Link: https://lore.kernel.org/r/20230724145000.125880-1-robdclark@gmail.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b88563ff21d429a4feb123f03afd3e0fe8a805b6
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jul 24 12:49:22 2023 +0200

    interconnect: qcom: qcm2290: Enable sync state

    [ Upstream commit 4e048e9b7a160f7112069c0ec2947be15f3e8154 ]

    Enable the generic .sync_state callback to ensure there are no
    outstanding votes that would waste power.

    Generally one would need a bunch of interface clocks to access the QoS
    registers when trying to go over all possible nodes during sync_state,
    but QCM2290 surprisingly does not seem to require any such handling.

    Fixes: 1a14b1ac3935 ("interconnect: qcom: Add QCM2290 driver support")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230720-topic-qcm2290_icc-v2-2-a2ceb9d3e713@linaro.org
    Signed-off-by: Georgi Djakov <djakov@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 20c83132a6a1d42dd45ed80cba9b4c80225dd000
Author: Ruidong Tian <tianruidong@linux.alibaba.com>
Date:   Fri Aug 4 16:15:14 2023 +0800

    coresight: tmc: Explicit type conversions to prevent integer overflow

    [ Upstream commit fd380097cdb305582b7a1f9476391330299d2c59 ]

    Perf cs_etm session executed unexpectedly when AUX buffer > 1G.

      perf record -C 0 -m ,2G -e cs_etm// -- <workload>
      [ perf record: Captured and wrote 2.615 MB perf.data ]

    Perf only collect about 2M perf data rather than 2G. This is becasuse
    the operation, "nr_pages << PAGE_SHIFT", in coresight tmc driver, will
    overflow when nr_pages >= 0x80000(correspond to 1G AUX buffer). The
    overflow cause buffer allocation to fail, and TMC driver will alloc
    minimal buffer size(1M). You can just get about 2M perf data(1M AUX
    buffer + perf data header) at least.

    Explicit convert nr_pages to 64 bit to avoid overflow.

    Fixes: 22f429f19c41 ("coresight: etm-perf: Add support for ETR backend")
    Fixes: 99443ea19e8b ("coresight: Add generic TMC sg table framework")
    Fixes: 2e499bbc1a92 ("coresight: tmc: implementing TMC-ETF AUX space API")
    Signed-off-by: Ruidong Tian <tianruidong@linux.alibaba.com>
    Reviewed-by: James Clark <james.clark@arm.com>
    Signed-off-by: Suzuki K Poulose <suzuki.poulose@arm.com>
    Link: https://lore.kernel.org/r/20230804081514.120171-2-tianruidong@linux.alibaba.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 327e8977f6b048db7ade42775a44734d4cc72478
Author: Gustavo A. R. Silva <gustavoars@kernel.org>
Date:   Wed Aug 2 08:46:26 2023 -0600

    RDMA/irdma: Replace one-element array with flexible-array member

    [ Upstream commit 38313c6d2a02c28162e06753b01bd885caf9386d ]

    One-element and zero-length arrays are deprecated. So, replace
    one-element array in struct irdma_qvlist_info with flexible-array
    member.

    A patch for this was sent a while ago[1]. However, it seems that, at
    the time, the changes were partially folded[2][3], and the actual
    flexible-array transformation was omitted. This patch fixes that.

    The only binary difference seen before/after changes is shown below:

    |  drivers/infiniband/hw/irdma/hw.o
    | @@ -868,7 +868,7 @@
    | drivers/infiniband/hw/irdma/hw.c:484 (discriminator 2)
    |       size += struct_size(iw_qvlist, qv_info, rf->msix_count);
    |      55b:      imul   $0x45c,%rdi,%rdi
    |-     562:      add    $0x10,%rdi
    |+     562:      add    $0x4,%rdi

    which is, of course, expected as it reflects the mistake made
    while folding the patch I've mentioned above.

    Worth mentioning is the fact that with this change we save 12 bytes
    of memory, as can be inferred from the diff snapshot above. Notice
    that:

    $ pahole -C rdma_qv_info idrivers/infiniband/hw/irdma/hw.o
    struct irdma_qv_info {
            u32                     v_idx;               /*     0     4 */
            u16                     ceq_idx;             /*     4     2 */
            u16                     aeq_idx;             /*     6     2 */
            u8                      itr_idx;             /*     8     1 */

            /* size: 12, cachelines: 1, members: 4 */
            /* padding: 3 */
            /* last cacheline: 12 bytes */
    };

    Link: https://lore.kernel.org/linux-hardening/20210525230038.GA175516@embeddedor/ [1]
    Link: https://lore.kernel.org/linux-
hardening/bf46b428deef4e9e89b0ea1704b1f0e5@intel.com/ [2]
    Link: https://lore.kernel.org/linux-rdma/20210520143809.819-1-
shiraz.saleem@intel.com/T/#u [3]
    Fixes: 44d9e52977a1 ("RDMA/irdma: Implement device initialization definitions")
    Signed-off-by: Gustavo A. R. Silva <gustavoars@kernel.org>
    Link: https://lore.kernel.org/r/ZMpsQrZadBaJGkt4@work
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8b85e01dbfa8a8808f998af358314a2e7844fda0
Author: Oleksandr Natalenko <oleksandr@natalenko.name>
Date:   Mon Jul 31 10:40:34 2023 +0200

    scsi: qedf: Do not touch __user pointer in qedf_dbg_fp_int_cmd_read() directly

    [ Upstream commit 25dbc20deab5165f847b4eb42f376f725a986ee8 ]

    The qedf_dbg_fp_int_cmd_read() function invokes sprintf() directly on a
```

```
    __user pointer, which may crash the kernel.

    Avoid doing that by vmalloc()'ating a buffer for scnprintf() and then
    calling simple_read_from_buffer() which does a proper copy_to_user() call.

    Fixes: 61d8658b4a43 ("scsi: qedf: Add QLogic FastLinQ offload FCoE driver framework.")
    Link: https://lore.kernel.org/lkml/20230724120241.40495-1-oleksandr@redhat.com/
    Link: https://lore.kernel.org/linux-scsi/20230726101236.11922-1-skashyap@marvell.com/
    Cc: Saurav Kashyap <skashyap@marvell.com>
    Cc: Rob Evers <revers@redhat.com>
    Cc: Johannes Thumshirn <Johannes.Thumshirn@wdc.com>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: Jozef Bacik <jobacik@redhat.com>
    Cc: Laurence Oberman <loberman@redhat.com>
    Cc: "James E.J. Bottomley" <jejb@linux.ibm.com>
    Cc: "Martin K. Petersen" <martin.petersen@oracle.com>
    Cc: GR-QLogic-Storage-Upstream@marvell.com
    Cc: linux-scsi@vger.kernel.org
    Reviewed-by: Laurence Oberman <loberman@redhat.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Tested-by: Laurence Oberman <loberman@redhat.com>
    Acked-by: Saurav Kashyap <skashyap@marvell.com>
    Signed-off-by: Oleksandr Natalenko <oleksandr@redhat.com>
    Link: https://lore.kernel.org/r/20230731084034.37021-4-oleksandr@redhat.com
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ad087c80cd065760f8074a519f0f4f089f81c20
Author: Oleksandr Natalenko <oleksandr@natalenko.name>
Date:   Mon Jul 31 10:40:33 2023 +0200

    scsi: qedf: Do not touch __user pointer in qedf_dbg_debug_cmd_read() directly

    [ Upstream commit 31b5991a9a91ba97237ac9da509d78eec453ff72 ]

    The qedf_dbg_debug_cmd_read() function invokes sprintf() directly on a
    __user pointer, which may crash the kernel.

    Avoid doing that by using a small on-stack buffer for scnprintf() and then
    calling simple_read_from_buffer() which does a proper copy_to_user() call.

    Fixes: 61d8658b4a43 ("scsi: qedf: Add QLogic FastLinQ offload FCoE driver framework.")
    Link: https://lore.kernel.org/lkml/20230724120241.40495-1-oleksandr@redhat.com/
    Link: https://lore.kernel.org/linux-scsi/20230726101236.11922-1-skashyap@marvell.com/
    Cc: Saurav Kashyap <skashyap@marvell.com>
    Cc: Rob Evers <revers@redhat.com>
    Cc: Johannes Thumshirn <Johannes.Thumshirn@wdc.com>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: Jozef Bacik <jobacik@redhat.com>
    Cc: Laurence Oberman <loberman@redhat.com>
    Cc: "James E.J. Bottomley" <jejb@linux.ibm.com>
    Cc: "Martin K. Petersen" <martin.petersen@oracle.com>
    Cc: GR-QLogic-Storage-Upstream@marvell.com
    Cc: linux-scsi@vger.kernel.org
    Reviewed-by: Laurence Oberman <loberman@redhat.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Tested-by: Laurence Oberman <loberman@redhat.com>
    Acked-by: Saurav Kashyap <skashyap@marvell.com>
    Signed-off-by: Oleksandr Natalenko <oleksandr@redhat.com>
    Link: https://lore.kernel.org/r/20230731084034.37021-3-oleksandr@redhat.com
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7030c1643b9ba8d0ba96742925d2f2114c160c20
Author: Oleksandr Natalenko <oleksandr@natalenko.name>
```

```
Date:   Mon Jul 31 10:40:32 2023 +0200

    scsi: qedf: Do not touch __user pointer in qedf_dbg_stop_io_on_error_cmd_read() directly

    [ Upstream commit 7d3d20dee4f648ec44e9717d5f647d594d184433 ]

    The qedf_dbg_stop_io_on_error_cmd_read() function invokes sprintf()
    directly on a __user pointer, which may crash the kernel.

    Avoid doing that by using a small on-stack buffer for scnprintf() and then
    calling simple_read_from_buffer() which does a proper copy_to_user() call.

    Fixes: 61d8658b4a43 ("scsi: qedf: Add QLogic FastLinQ offload FCoE driver framework.")
    Link: https://lore.kernel.org/lkml/20230724120241.40495-1-oleksandr@redhat.com/
    Link: https://lore.kernel.org/linux-scsi/20230726101236.11922-1-skashyap@marvell.com/
    Cc: Saurav Kashyap <skashyap@marvell.com>
    Cc: Rob Evers <revers@redhat.com>
    Cc: Johannes Thumshirn <Johannes.Thumshirn@wdc.com>
    Cc: David Laight <David.Laight@ACULAB.COM>
    Cc: Jozef Bacik <jobacik@redhat.com>
    Cc: Laurence Oberman <loberman@redhat.com>
    Cc: "James E.J. Bottomley" <jejb@linux.ibm.com>
    Cc: "Martin K. Petersen" <martin.petersen@oracle.com>
    Cc: GR-QLogic-Storage-Upstream@marvell.com
    Cc: linux-scsi@vger.kernel.org
    Reviewed-by: Laurence Oberman <loberman@redhat.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Tested-by: Laurence Oberman <loberman@redhat.com>
    Acked-by: Saurav Kashyap <skashyap@marvell.com>
    Signed-off-by: Oleksandr Natalenko <oleksandr@redhat.com>
    Link: https://lore.kernel.org/r/20230731084034.37021-2-oleksandr@redhat.com
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 255c0e60e1d16874fc151358d94bc8df661600dd
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Fri Jul 21 15:07:49 2023 -0500

    RDMA/rxe: Fix incomplete state save in rxe_requester

    [ Upstream commit 5d122db2ff80cd2aed4dcd630befb56b51ddf947 ]

    If a send packet is dropped by the IP layer in rxe_requester()
    the call to rxe_xmit_packet() can fail with err == -EAGAIN.
    To recover, the state of the wqe is restored to the state before
    the packet was sent so it can be resent. However, the routines
    that save and restore the state miss a significnt part of the
    variable state in the wqe, the dma struct which is used to process
    through the sge table. And, the state is not saved before the packet
    is built which modifies the dma struct.

    Under heavy stress testing with many QPs on a fast node sending
    large messages to a slow node dropped packets are observed and
    the resent packets are corrupted because the dma struct was not
    restored. This patch fixes this behavior and allows the test cases
    to succeed.

    Fixes: 3050b9985024 ("IB/rxe: Fix race condition between requester and completer")
    Link: https://lore.kernel.org/r/20230721200748.4604-1-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c40dafb951a39f987131a3093b4f67c2231ca056
Author: Bob Pearson <rpearsonhpe@gmail.com>
```

```
Date:   Tue Jun 20 09:01:43 2023 -0500

    RDMA/rxe: Fix rxe_modify_srq

    [ Upstream commit cc28f351155def8db209647f2e20a59a7080825b ]

    This patch corrects an error in rxe_modify_srq where if the
    caller changes the srq size the actual new value is not returned
    to the caller since it may be larger than what is requested.
    Additionally it open codes the subroutine rcv_wqe_size() which
    adds very little value, and makes some whitespace changes.

    Fixes: 8700e3e7c485 ("Soft RoCE driver")
    Link: https://lore.kernel.org/r/20230620140142.9452-1-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d366642b3099bd322375f5b71ba84ab1d586cd6d
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Tue Jun 20 08:55:21 2023 -0500

    RDMA/rxe: Fix unsafe drain work queue code

    [ Upstream commit 5993b75d0bc71cd2b441d174b028fc36180f032c ]

    If create_qp does not fully succeed it is possible for qp cleanup
    code to attempt to drain the send or recv work queues before the
    queues have been created causing a seg fault. This patch checks
    to see if the queues exist before attempting to drain them.

    Link: https://lore.kernel.org/r/20230620135519.9365-3-rpearsonhpe@gmail.com
    Reported-by: syzbot+2da1965168e7dbcba136@syzkaller.appspotmail.com
    Closes: https://lore.kernel.org/linux-rdma/00000000000012d89205fe7cfe00@google.com/raw
    Fixes: 49dc9c1f0c7e ("RDMA/rxe: Cleanup reset state handling in rxe_resp.c")
    Fixes: fbdeb828a21f ("RDMA/rxe: Cleanup error state handling in rxe_comp.c")
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cbd45bbda937ec7be589c896980c2c3949e46b55
Author: Bob Pearson <rpearsonhpe@gmail.com>
Date:   Tue Jun 20 08:55:19 2023 -0500

    RDMA/rxe: Move work queue code to subroutines

    [ Upstream commit e0ba8ff46704fc924e2ef0451ba196cbdc0d68f2 ]

    This patch:
            - Moves code to initialize a qp send work queue to a
              subroutine named rxe_init_sq.
            - Moves code to initialize a qp recv work queue to a
              subroutine named rxe_init_rq.
            - Moves initialization of qp request and response packet
              queues ahead of work queue initialization so that cleanup
              of a qp if it is not fully completed can successfully
              attempt to drain the packet queues without a seg fault.
            - Makes minor whitespace cleanups.

    Fixes: 8700e3e7c485 ("Soft RoCE driver")
    Link: https://lore.kernel.org/r/20230620135519.9365-2-rpearsonhpe@gmail.com
    Signed-off-by: Bob Pearson <rpearsonhpe@gmail.com>
    Acked-by: Zhu Yanjun <zyjzyj2000@gmail.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit d978b28ba75fce1183b74a8ebfb38b08ffbf0351
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Thu Jul 27 18:11:20 2023 -0700

    x86/APM: drop the duplicate APM_MINOR_DEV macro

    [ Upstream commit 4ba2909638a29630a346d6c4907a3105409bee7d ]

    This source file already includes <linux/miscdevice.h>, which contains
    the same macro. It doesn't need to be defined here again.

    Fixes: 874bcd00f520 ("apm-emulation: move APM_MINOR_DEV to include/linux/miscdevice.h")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Jiri Kosina <jikos@kernel.org>
    Cc: x86@kernel.org
    Cc: Sohil Mehta <sohil.mehta@intel.com>
    Cc: Corentin Labbe <clabbe.montjoie@gmail.com>
    Reviewed-by: Sohil Mehta <sohil.mehta@intel.com>
    Link: https://lore.kernel.org/r/20230728011120.759-1-rdunlap@infradead.org
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4ee715e54e255b1be65722f715fca939d5c2ca7a
Author: Chunyan Zhang <chunyan.zhang@unisoc.com>
Date:   Tue Jul 25 14:40:53 2023 +0800

    serial: sprd: Fix DMA buffer leak issue

    [ Upstream commit cd119fdc3ee1450fbf7f78862b5de44c42b6e47f ]

    Release DMA buffer when _probe() returns failure to avoid memory leak.

    Fixes: f4487db58eb7 ("serial: sprd: Add DMA mode support")
    Signed-off-by: Chunyan Zhang <chunyan.zhang@unisoc.com>
    Reviewed-by: Baolin Wang <baolin.wang@linux.alibaba.com>
    Link: https://lore.kernel.org/r/20230725064053.235448-2-chunyan.zhang@unisoc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f5330832f979f437e5affe62059da71dc9d85418
Author: Chunyan Zhang <chunyan.zhang@unisoc.com>
Date:   Tue Jul 25 14:40:52 2023 +0800

    serial: sprd: Assign sprd_port after initialized to avoid wrong access

    [ Upstream commit f9608f1887568b728839d006024585ab02ef29e5 ]

    The global pointer 'sprd_port' may not zero when sprd_probe returns
    failure, that is a risk for sprd_port to be accessed afterward, and
    may lead to unexpected errors.

    For example:

    There are two UART ports, UART1 is used for console and configured in
    kernel command line, i.e. "console=";

    The UART1 probe failed and the memory allocated to sprd_port[1] was
    released, but sprd_port[1] was not set to NULL;

    In UART2 probe, the same virtual address was allocated to sprd_port[2],
    and UART2 probe process finally will go into sprd_console_setup() to
    register UART1 as console since it is configured as preferred console
    (filled to console_cmdline[]), but the console parameters (sprd_port[1])
    belong to UART2.
```

```
    So move the sprd_port[] assignment to where the port already initialized
    can avoid the above issue.

    Fixes: b7396a38fb28 ("tty/serial: Add Spreadtrum sc9836-uart driver support")
    Signed-off-by: Chunyan Zhang <chunyan.zhang@unisoc.com>
    Link: https://lore.kernel.org/r/20230725064053.235448-1-chunyan.zhang@unisoc.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e70e88098ecd4a2496f580db20c8c2d8219b564
Author: Biju Das <biju.das.jz@bp.renesas.com>
Date:   Tue Jul 25 18:16:23 2023 +0100

    iio: accel: adxl313: Fix adxl313_i2c_id[] table

    [ Upstream commit f636554c4cd1c644109cc525900a056495b86cc9 ]

    The .driver_data in adxl313_i2c_id[] for adxl312 and adxl314 is
    wrong. Fix this issue by adding corresponding adxl31x_chip_info
    data.

    Reported-by: Jonathan Cameron <jic23@kernel.org>
    Closes: https://lore.kernel.org/all/20230722172832.04ad7738@jic23-huawei
    Fixes: a7a1c60bc4c9 ("drivers: iio: accel: adxl312 and adxl314 support")
    Signed-off-by: Biju Das <biju.das.jz@bp.renesas.com>
    Reviewed-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link: https://lore.kernel.org/r/20230725171624.331283-2-biju.das.jz@bp.renesas.com
    Signed-off-by: Jonathan Cameron <Jonathan.Cameron@huawei.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b018c0440b871d8b001c996e95fa4538bd292de6
Author: Lin Ma <linma@zju.edu.cn>
Date:   Sun Jul 23 16:00:53 2023 +0800

    scsi: qla4xxx: Add length check when parsing nlattrs

    [ Upstream commit 47cd3770e31df942e2bb925a9a855c79ed0662eb ]

    There are three places that qla4xxx parses nlattrs:

     - qla4xxx_set_chap_entry()

     - qla4xxx_iface_set_param()

     - qla4xxx_sysfs_ddb_set_param()

    and each of them directly converts the nlattr to specific pointer of
    structure without length checking. This could be dangerous as those
    attributes are not validated and a malformed nlattr (e.g., length 0) could
    result in an OOB read that leaks heap dirty data.

    Add the nla_len check before accessing the nlattr data and return EINVAL if
    the length check fails.

    Fixes: 26ffd7b45fe9 ("[SCSI] qla4xxx: Add support to set CHAP entries")
    Fixes: 1e9e2be3ee03 ("[SCSI] qla4xxx: Add flash node mgmt support")
    Fixes: 00c31889f751 ("[SCSI] qla4xxx: fix data alignment and use nl helpers")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
    Link: https://lore.kernel.org/r/20230723080053.3714534-1-linma@zju.edu.cn
    Reviewed-by: Chris Leech <cleech@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a19e5b0cfbb7acace582ff7130e92de2e96112f
```

```
Author: Lin Ma <linma@zju.edu.cn>
Date:   Sun Jul 23 15:59:38 2023 +0800

    scsi: be2iscsi: Add length check when parsing nlattrs

    [ Upstream commit ee0268f230f66cb472df3424f380ea668da2749a ]

    beiscsi_iface_set_param() parses nlattr with nla_for_each_attr and assumes
    every attributes can be viewed as struct iscsi_iface_param_info.

    This is not true because there is no any nla_policy to validate the
    attributes passed from the upper function iscsi_set_iface_params().

    Add the nla_len check before accessing the nlattr data and return EINVAL if
    the length check fails.

    Fixes: 0e43895ec1f4 ("[SCSI] be2iscsi: adding functionality to change network settings
using iscsiadm")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
    Link: https://lore.kernel.org/r/20230723075938.3713864-1-linma@zju.edu.cn
    Reviewed-by: Chris Leech <cleech@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 665acfe7248497da14f6ddfdc0229efaf7cbb5f0
Author: Lin Ma <linma@zju.edu.cn>
Date:   Sun Jul 23 15:58:20 2023 +0800

    scsi: iscsi: Add strlen() check in iscsi_if_set{_host}_param()

    [ Upstream commit ce51c817008450ef4188471db31639d42d37a5e1 ]

    The functions iscsi_if_set_param() and iscsi_if_set_host_param() convert an
    nlattr payload to type char* and then call C string handling functions like
    sscanf and kstrdup:

      char *data = (char*)ev + sizeof(*ev);
      ...
      sscanf(data, "%d", &value);

    However, since the nlattr is provided by the user-space program and the
    nlmsg skb is allocated with GFP_KERNEL instead of GFP_ZERO flag (see
    netlink_alloc_large_skb() in netlink_sendmsg()), dirty data on the heap can
    lead to an OOB access for those string handling functions.

    By investigating how the bug is introduced, we find it is really
    interesting as the old version parsing code starting from commit
    fd7255f51a13 ("[SCSI] iscsi: add sysfs attrs for uspace sync up") treated
    the nlattr as integer bytes instead of string and had length check in
    iscsi_copy_param():

      if (ev->u.set_param.len != sizeof(uint32_t))
        BUG();

    But, since the commit a54a52caad4b ("[SCSI] iscsi: fixup set/get param
    functions"), the code treated the nlattr as C string while forgetting to
    add any strlen checks(), opening the possibility of an OOB access.

    Fix the potential OOB by adding the strlen() check before accessing the
    buf. If the data passes this check, all low-level set_param handlers can
    safely treat this buf as legal C string.

    Fixes: fd7255f51a13 ("[SCSI] iscsi: add sysfs attrs for uspace sync up")
    Fixes: 1d9bf13a9cf9 ("[SCSI] iscsi class: add iscsi host set param event")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
```

```
    Link: https://lore.kernel.org/r/20230723075820.3713119-1-linma@zju.edu.cn
    Reviewed-by: Chris Leech <cleech@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 698afeedff07e81e6453d69c3a12dadf7037a69f
Author: Lin Ma <linma@zju.edu.cn>
Date:   Tue Jul 25 10:45:29 2023 +0800

    scsi: iscsi: Add length check for nlattr payload

    [ Upstream commit 971dfcb74a800047952f5288512b9c7ddedb050a ]

    The current NETLINK_ISCSI netlink parsing loop checks every nlmsg to make
    sure the length is bigger than sizeof(struct iscsi_uevent) and then calls
    iscsi_if_recv_msg().

      nlh = nlmsg_hdr(skb);
      if (nlh->nlmsg_len < sizeof(*nlh) + sizeof(*ev) ||
        skb->len < nlh->nlmsg_len) {
        break;
      }
      ...
      err = iscsi_if_recv_msg(skb, nlh, &group);

    Hence, in iscsi_if_recv_msg() the nlmsg_data can be safely converted to
    iscsi_uevent as the length is already checked.

    However, in other cases the length of nlattr payload is not checked before
    the payload is converted to other data structures. One example is
    iscsi_set_path() which converts the payload to type iscsi_path without any
    checks:

      params = (struct iscsi_path *)((char *)ev + sizeof(*ev));

    Whereas iscsi_if_transport_conn() correctly checks the pdu_len:

      pdu_len = nlh->nlmsg_len - sizeof(*nlh) - sizeof(*ev);
      if ((ev->u.send_pdu.hdr_size > pdu_len) ..
        err = -EINVAL;

    To sum up, some code paths called in iscsi_if_recv_msg() do not check the
    length of the data (see below picture) and directly convert the data to
    another data structure. This could result in an out-of-bound reads and heap
    dirty data leakage.

                 _________  nlmsg_len(nlh) _________________
                /                                           \
    +----------+--------------+---------------------------+
    | nlmsghdr | iscsi_uevent |          data             |
    +----------+--------------+---------------------------+
                             \                            /
                              iscsi_uevent->u.set_param.len

    Fix the issue by adding the length check before accessing it. To clean up
    the code, an additional parameter named rlen is added. The rlen is
    calculated at the beginning of iscsi_if_recv_msg() which avoids duplicated
    calculation.

    Fixes: ac20c7bf070d ("[SCSI] iscsi_transport: Added Ping support")
    Fixes: 43514774ff40 ("[SCSI] iscsi class: Add new NETLINK_ISCSI messages for cnic/bnx2i
driver.")
    Fixes: 1d9bf13a9cf9 ("[SCSI] iscsi class: add iscsi host set param event")
    Fixes: 01cb225dad8d ("[SCSI] iscsi: add target discvery event to transport class")
    Fixes: 264faaaa1254 ("[SCSI] iscsi: add transport end point callbacks")
```

```
    Fixes: fd7255f51a13 ("[SCSI] iscsi: add sysfs attrs for uspace sync up")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
    Link: https://lore.kernel.org/r/20230725024529.428311-1-linma@zju.edu.cn
    Reviewed-by: Chris Leech <cleech@redhat.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b648bd321b2f4a4ed46415658cddbea19d943200
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Mon Jul 24 13:08:29 2023 -0700

    scsi: ufs: Fix residual handling

    [ Upstream commit 2903265e27bfc6dea915dd9e17a1b2587f621f73 ]

    Only call scsi_set_resid() in case of an underflow. Do not call
    scsi_set_resid() in case of an overflow.

    Cc: Avri Altman <avri.altman@wdc.com>
    Cc: Adrian Hunter <adrian.hunter@intel.com>
    Fixes: cb38845d90fc ("scsi: ufs: core: Set the residual byte count")
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Link: https://lore.kernel.org/r/20230724200843.3376570-2-bvanassche@acm.org
    Reviewed-by: Avri Altman <avri.altman@wdc.com>
    Reviewed-by: Adrian Hunter <adrian.hunter@intel.com>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9533f0397322b0050f496698705df5ea6a867e5e
Author: Bart Van Assche <bvanassche@acm.org>
Date:   Mon Jul 24 13:08:30 2023 -0700

    scsi: RDMA/srp: Fix residual handling

    [ Upstream commit 89e637c19b2441aabc8dbf22a8745b932fd6996e ]

    Although the code for residual handling in the SRP initiator follows the
    SCSI documentation, that documentation has never been correct. Because
    scsi_finish_command() starts from the data buffer length and subtracts the
    residual, scsi_set_resid() must not be called if a residual overflow
    occurs. Hence remove the scsi_set_resid() calls from the SRP initiator if a
    residual overflow occurrs.

    Cc: Leon Romanovsky <leon@kernel.org>
    Cc: Jason Gunthorpe <jgg@nvidia.com>
    Fixes: 9237f04e12cc ("scsi: core: Fix scsi_get/set_resid() interface")
    Fixes: e714531a349f ("IB/srp: Fix residual handling")
    Signed-off-by: Bart Van Assche <bvanassche@acm.org>
    Link: https://lore.kernel.org/r/20230724200843.3376570-3-bvanassche@acm.org
    Acked-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0fe57f9e9400c850b4eae724e134e4601384087b
Author: Xu Yang <xu.yang_2@nxp.com>
Date:   Tue Jun 27 19:03:52 2023 +0800

    usb: phy: mxs: fix getting wrong state with mxs_phy_is_otg_host()

    [ Upstream commit 5eda42aebb7668b4dcff025cd3ccb0d3d7c53da6 ]

    The function mxs_phy_is_otg_host() will return true if OTG_ID_VALUE is
    0 at USBPHY_CTRL register. However, OTG_ID_VALUE will not reflect the real
    state if the ID pin is float, such as Host-only or Type-C cases. The value
    of OTG_ID_VALUE is always 1 which means device mode.
```

```
    This patch will fix the issue by judging the current mode based on
    last_event. The controller will update last_event in time.

    Fixes: 7b09e67639d6 ("usb: phy: mxs: refine mxs_phy_disconnect_line")
    Signed-off-by: Xu Yang <xu.yang_2@nxp.com>
    Acked-by: Peter Chen <peter.chen@kernel.org>
    Link: https://lore.kernel.org/r/20230627110353.1879477-2-xu.yang_2@nxp.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 451dc187cadd47771e5d9434fe220fad7be84057
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Wed Jun 14 16:06:47 2023 +0300

    media: mediatek: vcodec: fix resource leaks in vdec_msg_queue_init()

    [ Upstream commit cf10b0bb503c974ba049d6f888b21178be20a962 ]

    If we encounter any error in the vdec_msg_queue_init() then we need
    to set "msg_queue->wdma_addr.size = 0;".  Normally, this is done
    inside the vdec_msg_queue_deinit() function.  However, if the
    first call to allocate &msg_queue->wdma_addr fails, then the
    vdec_msg_queue_deinit() function is a no-op.  For that situation, just
    set the size to zero explicitly and return.

    There were two other error paths which did not clean up before returning.
    Change those error paths to goto mem_alloc_err.

    Fixes: b199fe46f35c ("media: mtk-vcodec: Add msg queue feature for lat and core
architecture")
    Fixes: 2f5d0aef37c6 ("media: mediatek: vcodec: support stateless AV1 decoder")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33940d812d9f556609646846a6b93d04880d654a
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Wed Jun 14 16:05:39 2023 +0300

    media: mediatek: vcodec: fix potential double free

    [ Upstream commit be40f524b6edac4fb9a98ef79620fd9b9497a998 ]

    The "lat_buf->private_data" needs to be set to NULL to prevent a
    double free.  How this would happen is if vdec_msg_queue_init() failed
    twice in a row and on the second time it failed earlier than on the
    first time.

    The vdec_msg_queue_init() function has a loop which does:
            for (i = 0; i < NUM_BUFFER_COUNT; i++) {

    Each iteration initializes one element in the msg_queue->lat_buf[] array
    and then the clean up function vdec_msg_queue_deinit() frees each
    element of the msg_queue->lat_buf[] array.  This clean up code relies
    on the assumption that every element is either initialized or zeroed.
    Leaving a freed pointer which is non-zero breaks the assumption.

    Fixes: b199fe46f35c ("media: mtk-vcodec: Add msg queue feature for lat and core
architecture")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit a99c0a804b2cb71a7100b734908dfc3874a5906b
Author: Irui Wang <irui.wang@mediatek.com>
Date:   Wed Jul 5 17:14:41 2023 +0800

    media: mediatek: vcodec: Return NULL if no vdec_fb is found

    [ Upstream commit dfa2d6e07432270330ae191f50a0e70636a4cd2b ]

    "fb_use_list" is used to store used or referenced frame buffers for
    vp9 stateful decoder. "NULL" should be returned when getting target
    frame buffer failed from "fb_use_list", not a random unexpected one.

    Fixes: f77e89854b3e ("[media] vcodec: mediatek: Add Mediatek VP9 Video Decoder Driver")
    Signed-off-by: Irui Wang <irui.wang@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e298348d978ef20d8b3a73e946918e8dcced4925
Author: Xiaoyong Lu <xiaoyong.lu@mediatek.com>
Date:   Tue Jul 4 09:51:35 2023 +0800

    media: mediatek: vcodec: fix AV1 decode fail for 36bit iova

    [ Upstream commit 89a4f369b20810a8365f87badf7862c67d344bbe ]

    Fix av1 decode fail when iova is 36bit.

    Decoder hardware will access incorrect iova address when tile buffer is
    36bit, it will lead to iommu fault when hardware access dram data.

    Fixes: 2f5d0aef37c6 ("media: mediatek: vcodec: support stateless AV1 decoder")
    Signed-off-by: Xiaoyong Lu<xiaoyong.lu@mediatek.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 882ed5bd39e82bc1b620a4d2f5a461436eff5df2
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jul 18 17:50:13 2023 +0800

    media: amphion: ensure the bitops don't cross boundaries

    [ Upstream commit 5bd28eae48589694ff4e5badb03bf75dae695b3f ]

    the supported_instance_count determine the instance index range,
    it shouldn't exceed the bits number of instance_mask,
    otherwise the bitops of instance_mask may cross boundaries

    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3056d01a185c553ff35c02851d6d119a6c7bb6c5
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jul 18 17:50:12 2023 +0800

    media: amphion: fix UNUSED_VALUE issue reported by coverity

    [ Upstream commit cf6a06354989c41b536be8e094561ee16223cf1f ]

    assign value '-EINVAL' to ret, but the stored value is overwritten
    before it can be used
```

```
    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bc218569b341012db6e6a92e4fa5bba3806e25f
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jul 18 17:50:11 2023 +0800

    media: amphion: fix UNINIT issues reported by coverity

    [ Upstream commit c224d0497a31ea2d173e1ea16af308945bff9037 ]

    using uninitialized value may introduce risk

    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ef56b2db216f130c4240aed907d1c5272c2d298d
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jul 18 17:50:10 2023 +0800

    media: amphion: fix REVERSE_INULL issues reported by coverity

    [ Upstream commit 79d3bafaecc13bccab1ebbd28a15e669c5a4cdaf ]

    null-checking of a pointor is suggested before dereferencing it

    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90e632cb944c57c657788d7f8ea97a3680e2afa1
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jul 18 17:50:09 2023 +0800

    media: amphion: fix CHECKED_RETURN issues reported by coverity

    [ Upstream commit b237b058adbc7825da9c8f358f1ff3f0467d623a ]

    calling "vpu_cmd_send/vpu_get_buffer_state/vpu_session_alloc_fs"
    without checking return value

    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9b3a491a0c8386f86da41e190b27d81454766f70
Author: Benjamin Gaignard <benjamin.gaignard@collabora.com>
Date:   Mon Jul 17 17:06:11 2023 +0200

    media: rkvdec: increase max supported height for H.264

    [ Upstream commit f000e6ca2d60fefd02a180a57df2c4162fa0c1b7 ]

    After testing it is possible for the hardware to decode H264
    bistream with a height up to 2560.
```

```
    Signed-off-by: Benjamin Gaignard <benjamin.gaignard@collabora.com>
    Fixes: cd33c830448ba ("media: rkvdec: Add the rkvdec driver")
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9090e747458c62ded7623436dcd6f10feab54ad
Author: Ming Qian <ming.qian@nxp.com>
Date:   Mon Jul 10 15:44:11 2023 +0800

    media: amphion: decoder support display delay for all formats

    [ Upstream commit b69713f502027150ecc08e663fa1804d78b3ef42 ]

    the firmware only support low latency mode for h264,
    but firmware will notify an event to driver
    when one frame is decoded,
    if V4L2_CID_MPEG_VIDEO_DEC_DISPLAY_DELAY_ENABLE is enabled,
    and V4L2_CID_MPEG_VIDEO_DEC_DISPLAY_DELAY is set to 0,
    driver can display the decoded frame immediately.

    Fixes: ffa331d9bf94 ("media: amphion: decoder implement display delay enable")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2fc20f8bcc2b4d31c808a5320506c31aa2cf3834
Author: Zheng Wang <zyytlz.wz@163.com>
Date:   Fri Jul 7 17:24:14 2023 +0800

    media: mtk-jpeg: Fix use after free bug due to uncanceled work

    [ Upstream commit c677d7ae83141d390d1253abebafa49c962afb52 ]

    In mtk_jpeg_probe, &jpeg->job_timeout_work is bound with
    mtk_jpeg_job_timeout_work. Then mtk_jpeg_dec_device_run
    and mtk_jpeg_enc_device_run may be called to start the
    work.
    If we remove the module which will call mtk_jpeg_remove
    to make cleanup, there may be a unfinished work. The
    possible sequence is as follows, which will cause a
    typical UAF bug.

    Fix it by canceling the work before cleanup in the mtk_jpeg_remove

    CPU0                  CPU1

                         |mtk_jpeg_job_timeout_work
    mtk_jpeg_remove      |
      v4l2_m2m_release   |
        kfree(m2m_dev);  |
                         |
                         | v4l2_m2m_get_curr_priv
                         |   m2m_dev->curr_ctx //use
    Fixes: b2f0d2724ba4 ("[media] vcodec: mediatek: Add Mediatek JPEG Decoder Driver")
    Signed-off-by: Zheng Wang <zyytlz.wz@163.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Reviewed-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 07750bb9c5b7ff0edbc70f858c9380fdd4aaaa44
Author: Michael Tretter <m.tretter@pengutronix.de>
```

```
Date:   Thu Jul 6 09:15:10 2023 +0200

    media: verisilicon: Fix TRY_FMT on encoder OUTPUT

    [ Upstream commit b3b4c9d3cb3bf8725a3ded26f7042b1a37f25333 ]

    Commit f100ce3bbd6a ("media: verisilicon: Fix crash when probing
    encoder") removed vpu_fmt from hantro_try_fmt(), since it was
    initialized from vpu_dst_fmt, which may not be initialized, when TRY_FMT
    is called. It was replaced by fmt, which is found using the pixelformat.

    For the encoder, this changed the fmt to contain the raw format instead
    of the coded format. The format constraints as of fmt->frmsize are only
    valid for the coded format and are 0 for the raw formats. Therefore, the
    size of a encoder OUTPUT device is constrained to 0 and the
    v4l2-compliance tests for G_FMT, TRY_FMT, and SET_FMT fail.

    Bring back vpu_fmt to use the coded format on an encoder OUTPUT device,
    but initialize it using the currently set pixelformat on dst_fmt, which
    is the coded format on an encoder.

    Fixes: f100ce3bbd6a ("media: verisilicon: Fix crash when probing encoder")
    Signed-off-by: Michael Tretter <m.tretter@pengutronix.de>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9284217dcc2ec3396a10d6b9ca16806cb65eb954
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jun 13 17:14:08 2023 +0800

    media: amphion: add helper function to get id name

    [ Upstream commit 12cd8b8ac02525977b2e860a877add10e8ce7468 ]

    convert numbers into meaningful names,
    then it can improve the log readability

    Fixes: 9f599f351e86 ("media: amphion: add vpu core driver")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d3973bcfbb6e0f307cf8e0eb794816d8992a7ed9
Author: Ming Qian <ming.qian@nxp.com>
Date:   Tue Jun 13 15:48:46 2023 +0800

    media: amphion: reinit vpu if reqbufs output 0

    [ Upstream commit 73e3f09292a0492a3fe0f87a8170a74f12624c5e ]

    according to v4l2 stateful decoder document 4.5.1.3. State Machine,
    the state should change from seek to initialization
    if call VIDIOC_REQBUFS(OUTPUT, 0).

    so reinit the vpu decoder if reqbufs output 0

    Fixes: 6de8d628df6e ("media: amphion: add v4l2 m2m vpu decoder stateful driver")
    Signed-off-by: Ming Qian <ming.qian@nxp.com>
    Tested-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Reviewed-by: Nicolas Dufresne <nicolas.dufresne@collabora.com>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 78fa7113f6dd8f6ca31cbbcaee99dc95c4c1bcfa
```

```
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 10:01:40 2023 +0200

    dt-bindings: extcon: maxim,max77843: restrict connector properties

    [ Upstream commit fb2c3f72e819254d8c76de95917e5f9ff232586c ]

    Do not allow any other properties in connector child, except what
    usb-connector.yaml evaluates.

    Fixes: 9729cad0278b ("dt-bindings: extcon: maxim,max77843: Add MAX77843 bindings")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Chanwoo Choi <cw00.choi@samsung.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf71d21b411fe7bb3bee12dc4b717d5004067fe5
Author: Xingui Yang <yangxingui@huawei.com>
Date:   Tue Jul 11 11:14:58 2023 +0800

    scsi: hisi_sas: Fix normally completed I/O analysed as failed

    [ Upstream commit f5393a5602cacfda2014e0ff8220e5a7564e7cd1 ]

    The PIO read command has no response frame and the struct iu[1024] won't be
    filled. I/Os which are normally completed will be treated as failed in
    sas_ata_task_done() when iu contains abnormal dirty data.

    Consequently ending_fis should not be filled by iu when the response frame
    hasn't been written to memory.

    Fixes: d380f55503ed ("scsi: hisi_sas: Don't bother clearing status buffer IU in task
prep")
    Signed-off-by: Xingui Yang <yangxingui@huawei.com>
    Signed-off-by: Xiang Chen <chenxiang66@hisilicon.com>
    Link: https://lore.kernel.org/r/1689045300-44318-2-git-send-email-
chenxiang66@hisilicon.com
    Signed-off-by: Martin K. Petersen <martin.petersen@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 95cf458f1652cf7c40cda243c36ed9120c6d84aa
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Mon Jul 17 11:12:12 2023 -0400

    RDMA/siw: Fabricate a GID on tun and loopback devices

    [ Upstream commit bad5b6e34ffbaacc77ad28a0f482e33b3929e635 ]

    LOOPBACK and NONE (tunnel) devices have all-zero MAC addresses.
    Currently, siw_device_create() falls back to copying the IB device's
    name in those cases, because an all-zero MAC address breaks the RDMA
    core address resolution mechanism.

    However, at the point when siw_device_create() constructs a GID, the
    ib_device::name field is uninitialized, leaving the MAC address to
    remain in an all-zero state.

    Fabricate a random artificial GID for such devices, and ensure this
    artificial GID is returned for all device query operations.

    Link:
https://lore.kernel.org/r/168960673260.3007.12378736853793339110.stgit@manet.1015granger.net
    Reported-by: Tom Talpey <tom@talpey.com>
    Fixes: a2d36b02c15d ("RDMA/siw: Enable siw on tunnel devices")
    Reviewed-by: Bernard Metzler <bmt@zurich.ibm.com>
    Reviewed-by: Tom Talpey <tom@talpey.com>
```

```
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92dba99007a98044326b828a84d4df93d90328c4
Author: Daniil Dulov <d.dulov@aladdin.ru>
Date:   Fri Jun 2 01:55:01 2023 -0700

    media: cx24120: Add retval check for cx24120_message_send()

    [ Upstream commit 96002c0ac824e1773d3f706b1f92e2a9f2988047 ]

    If cx24120_message_send() returns error, we should keep local struct
    unchanged.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 5afc9a25be8d ("[media] Add support for TechniSat Skystar S2")
    Signed-off-by: Daniil Dulov <d.dulov@aladdin.ru>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2cc9f11aeae2887a4db25c27323fc445f4b49e86
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Mon May 29 07:58:36 2023 +0200

    media: dvb-usb: m920x: Fix a potential memory leak in m920x_i2c_xfer()

    [ Upstream commit ea9ef6c2e001c5dc94bee35ebd1c8a98621cf7b8 ]

    'read' is freed when it is known to be NULL, but not when a read error
    occurs.

    Revert the logic to avoid a small leak, should a m920x_read() call fail.

    Fixes: a2ab06d7c4d6 ("media: m920x: don't use stack on USB reads")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3b685ded7c775052faac9a322ce3562e4c5ab6d
Author: Daniil Dulov <d.dulov@aladdin.ru>
Date:   Fri Mar 24 06:38:32 2023 -0700

    media: dib7000p: Fix potential division by zero

    [ Upstream commit a1db7b2c5533fc67e2681eb5efc921a67bc7d5b8 ]

    Variable loopdiv can be assigned 0, then it is used as a denominator,
    without checking it for 0.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 713d54a8bd81 ("[media] DiB7090: add support for the dib7090 based")
    Signed-off-by: Daniil Dulov <d.dulov@aladdin.ru>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    [hverkuil: (bw != NULL) -> bw]
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2287a78a2cbfcd510a386a8414a381929ea0e831
Author: Dongliang Mu <dzm91@hust.edu.cn>
Date:   Mon Feb 27 18:24:08 2023 +0800

    drivers: usb: smsusb: fix error handling code in smsusb_init_device
```

```
    [ Upstream commit b9c7141f384097fa4fa67d2f72e5731d628aef7c ]

    The previous commit 4b208f8b561f ("[media] siano: register media controller
    earlier")moves siano_media_device_register before smscore_register_device,
    and adds corresponding error handling code if smscore_register_device
    fails. However, it misses the following error handling code of
    smsusb_init_device.

    Fix this by moving error handling code at the end of smsusb_init_device
    and adding a goto statement in the following error handling parts.

    Fixes: 4b208f8b561f ("[media] siano: register media controller earlier")
    Signed-off-by: Dongliang Mu <dzm91@hust.edu.cn>
    Signed-off-by: Hans Verkuil <hverkuil-cisco@xs4all.nl>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d08f524ff12d180f359b70979c95b18a2d43b8b
Author: Jonas Karlman <jonas@kwiboo.se>
Date:   Sat Jun 17 18:25:45 2023 +0000

    iommu: rockchip: Fix directory table address encoding

    [ Upstream commit 6df63b7ebdaf5fcd75dceedf6967d0761e56eca1 ]

    The physical address to the directory table is currently encoded using
    the following bit layout for IOMMU v2.

     31:12 - Address bit 31:0
     11: 4 - Address bit 39:32

    This is also the bit layout used by the vendor kernel.

    However, testing has shown that addresses to the directory/page tables
    and memory pages are all encoded using the same bit layout.

    IOMMU v1:
     31:12 - Address bit 31:0

    IOMMU v2:
     31:12 - Address bit 31:0
     11: 8 - Address bit 35:32
      7: 4 - Address bit 39:36

    Change to use the mk_dtentries ops to encode the directory table address
    correctly. The value written to DTE_ADDR may include the valid bit set,
    a bit that is ignored and DTE_ADDR reg read it back as 0.

    This also update the bit layout comment for the page address and the
    number of nybbles that are read back for DTE_ADDR comment.

    These changes render the dte_addr_phys and dma_addr_dte ops unused and
    is removed.

    Fixes: 227014b33f62 ("iommu: rockchip: Add internal ops to handle variants")
    Fixes: c55356c534aa ("iommu: rockchip: Add support for iommu v2")
    Fixes: c987b65a574f ("iommu/rockchip: Fix physical address decoding")
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Reviewed-by: Robin Murphy <robin.murphy@arm.com>
    Link: https://lore.kernel.org/r/20230617182540.3091374-2-jonas@kwiboo.se
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 98d86bf32187db27946ca817c2467a5f2f7aa02f
Author: Daniel Marcovitch <dmarcovitch@nvidia.com>
Date:   Fri Jun 9 10:51:45 2023 +0000
```

```
    iommu/amd/iommu_v2: Fix pasid_state refcount dec hit 0 warning on pasid unbind

    [ Upstream commit 534103bcd52ca9c1fecbc70e717b4a538dc4ded8 ]

    When unbinding pasid - a race condition exists vs outstanding page faults.

    To prevent this, the pasid_state object contains a refcount.
        * set to 1 on pasid bind
        * incremented on each ppr notification start
        * decremented on each ppr notification done
        * decremented on pasid unbind

    Since refcount_dec assumes that refcount will never reach 0:
      the current implementation causes the following to be invoked on
      pasid unbind:
            REFCOUNT_WARN("decrement hit 0; leaking memory")

    Fix this issue by changing refcount_dec to refcount_dec_and_test
    to explicitly handle refcount=1.

    Fixes: 8bc54824da4e ("iommu/amd: Convert from atomic_t to refcount_t on pasid_state-
>count")
    Signed-off-by: Daniel Marcovitch <dmarcovitch@nvidia.com>
    Signed-off-by: Vasant Hegde <vasant.hegde@amd.com>
    Link: https://lore.kernel.org/r/20230609105146.7773-2-vasant.hegde@amd.com
    Signed-off-by: Joerg Roedel <jroedel@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8a1cd87bb9fa3149ee112ecb8058908dc9b520e
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Wed Jun 14 20:31:05 2023 +0200

    media: v4l2-core: Fix a potential resource leak in v4l2_fwnode_parse_link()

    [ Upstream commit d7b13edd4cb4bfa335b6008ab867ac28582d3e5c ]

    If fwnode_graph_get_remote_endpoint() fails, 'fwnode' is known to be NULL,
    so fwnode_handle_put() is a no-op.

    Release the reference taken from a previous fwnode_graph_get_port_parent()
    call instead.

    Also handle fwnode_graph_get_port_parent() failures.

    In order to fix these issues, add an error handling path to the function
    and the needed gotos.

    Fixes: ca50c197bd96 ("[media] v4l: fwnode: Support generic fwnode for parsing
standardised properties")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1ce31523f1ef9f519fc138375e3a9b526df167a5
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Thu Jun 15 12:30:30 2023 +0200

    media: i2c: tvp5150: check return value of devm_kasprintf()

    [ Upstream commit 26ce7054d804be73935b9268d6e0ecf2fbbc8aef ]

    devm_kasprintf() returns a pointer to dynamically allocated memory.
    Pointer could be NULL in case allocation fails. Check pointer validity.
```

```
    Identified with coccinelle (kmerr.cocci script).

    Fixes: 0556f1d580d4 ("media: tvp5150: add input source selection of_graph support")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Reviewed-by: Marco Felsch <m.felsch@pengutronix.de>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1ee15d057607eb49b0755d1d4401efe2f09a5db0
Author: Hans de Goede <hdegoede@redhat.com>
Date:   Sun Jun 18 20:17:40 2023 +0200

    media: ad5820: Drop unsupported ad5823 from i2c_ and of_device_id tables

    [ Upstream commit f126ff7e4024f6704e6ec0d4137037568708a3c7 ]

    The supported ad5820 and ad5821 VCMs both use a single 16 bit register
    which is written by sending 2 bytes with the data directly after sending
    the i2c-client address.

    The ad5823 OTOH has a more typical i2c / smbus device setup with multiple
    8 bit registers where the first byte send after the i2c-client address is
    the register address and the actual data only starts from the second byte
    after the i2c-client address.

    The ad5823 i2c_ and of_device_id-s was added at the same time as
    the ad5821 ids with as rationale:

    """
    Some camera modules also refer that AD5823 is a replacement of AD5820:
    https://download.kamami.com/p564094-OV8865_DS.pdf
    """

    The AD5823 may be an electrical and functional replacement of the AD5820,
    but from a software pov it is not compatible at all and it is going to
    need its own driver, drop its id from the ad5820 driver.

    Fixes: b8bf73136bae ("media: ad5820: Add support for ad5821 and ad5823")
    Cc: Pavel Machek <pavel@ucw.cz>
    Cc: Ricardo Ribalda Delgado <ribalda@kernel.org>
    Signed-off-by: Hans de Goede <hdegoede@redhat.com>
    Reviewed-by: Ricardo Ribalda <ribalda@chromium.org>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bb28a5ed13923fac695a4d30af7ffa817cdfd3ce
Author: Tommaso Merciai <tomm.merciai@gmail.com>
Date:   Tue Jun 13 10:07:34 2023 +0200

    media: i2c: imx290: drop format param from imx290_ctrl_update

    [ Upstream commit 9b4e0e7a570d222be5f5e0f914d3c4528eadeeb4 ]

    The format param actually is not used in imx290_ctrl_update
    function, let's drop this

    Fixes: bc35f9a21a55 ("media: i2c: imx290: Fix the pixel rate at 148.5Mpix/s")
    Signed-off-by: Tommaso Merciai <tomm.merciai@gmail.com>
    Reviewed-by: Dave Stevenson <dave.stevenson@raspberrypi.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit aec4f926c018072f48898a49a1993242b7c91b0a
Author: Guoniu.zhou <guoniu.zhou@nxp.com>
Date:   Mon Jun 12 04:43:40 2023 +0200

    media: ov5640: fix low resolution image abnormal issue

    [ Upstream commit a828002f38c5ee49d3f0c0e64c0f0caa1aec8dc2 ]

    OV5640 will output abnormal image data when work at low resolution
    (320x240, 176x144 and 160x120) after switching from high resolution,
    such as 1080P, the time interval between high and low switching must
    be less than 1000ms in order to OV5640 don't enter suspend state during
    the time.

    The reason is by 0x3824 value don't restore to initialize value when
    do resolution switching. In high resolution setting array, 0x3824 is
    set to 0x04, but low resolution setting array remove 0x3824 in commit
    db15c1957a2d ("media: ov5640: Remove duplicated mode settings"). So
    when do resolution switching from high to low, such as 1080P to 320x240,
    and the time interval is less than auto suspend delay time which means
    global initialize setting array will not be loaded, the output image
    data are abnormal. Hence move 0x3824 from ov5640_init_setting[] table
    to ov5640_setting_low_res[] table and also move 0x4407 0x460b, 0x460c
    to avoid same issue.

    Fixes: db15c1957a2d ("media: ov5640: Remove duplicated mode settings")
    Signed-off-by: Guoniu.zhou <guoniu.zhou@nxp.com>
    Reviewed-by: Jacopo Mondi <jacopo.mondi@ideasonboard.com>
    Signed-off-by: Sakari Ailus <sakari.ailus@linux.intel.com>
    Signed-off-by: Mauro Carvalho Chehab <mchehab@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b85ee12aef35ab4b24c57ef33fe6ec683c86069f
Author: Minjie Du <duminjie@vivo.com>
Date:   Thu Jul 6 10:27:03 2023 +0800

    RDMA/qedr: Remove a duplicate assignment in irdma_query_ah()

    [ Upstream commit 65e02e840847158c7ee48ca8e6e91062b0f78662 ]

    Delete a duplicate statement from this function implementation.

    Fixes: b48c24c2d710 ("RDMA/irdma: Implement device supported verb APIs")
    Signed-off-by: Minjie Du <duminjie@vivo.com>
    Acked-by: Alok Prasad <palok@marvell.com>
    Link: https://lore.kernel.org/r/20230706022704.1260-1-duminjie@vivo.com
    Signed-off-by: Leon Romanovsky <leon@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7910f8e53c24228a1f89ea85fa1bde3a3efc5cff
Author: Waiman Long <longman@redhat.com>
Date:   Tue Jun 27 10:35:00 2023 -0400

    cgroup/cpuset: Inherit parent's load balance state in v2

    [ Upstream commit c8c926200c55454101f072a4b16c9ff5b8c9e56f ]

    Since commit f28e22441f35 ("cgroup/cpuset: Add a new isolated
    cpus.partition type"), the CS_SCHED_LOAD_BALANCE bit of a v2 cpuset
    can be on or off. The child cpusets of a partition root must have the
    same setting as its parent or it may screw up the rebuilding of sched
    domains. Fix this problem by making sure the a child v2 cpuset will
    follows its parent cpuset load balance state unless the child cpuset
    is a new partition root itself.
```

```
    Fixes: f28e22441f35 ("cgroup/cpuset: Add a new isolated cpus.partition type")
    Signed-off-by: Waiman Long <longman@redhat.com>
    Signed-off-by: Tejun Heo <tj@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 012e887bb9eaa2c4175a96ca81e9ff457685cf65
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Wed Aug 30 14:31:31 2023 -0400

    pNFS: Fix assignment of xprtdata.cred

    [ Upstream commit c4a123d2e8c4dc91d581ee7d05c0cd51a0273fab ]

    The comma at the end of the line was leftover from an earlier refactor
    of the _nfs4_pnfs_v3_ds_connect() function. This is technically valid C,
    so the compilers didn't catch it, but if I'm understanding how it works
    correctly it assigns the return value of rpc_clnt_add_xprtr() to
    xprtdata.cred.

    Reported-by: Olga Kornievskaia <kolga@netapp.com>
    Fixes: a12f996d3413 ("NFSv4/pNFS: Use connections to a DS that are all of the same
protocol family")
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6cf45efe0419abc569a0a7b10f51b2483a3b7b1
Author: Olga Kornievskaia <kolga@netapp.com>
Date:   Thu Aug 24 16:43:53 2023 -0400

    NFSv4.2: fix handling of COPY ERR_OFFLOAD_NO_REQ

    [ Upstream commit 5690eed941ab7e33c3c3d6b850100cabf740f075 ]

    If the client sent a synchronous copy and the server replied with
    ERR_OFFLOAD_NO_REQ indicating that it wants an asynchronous
    copy instead, the client should retry with asynchronous copy.

    Fixes: 539f57b3e0fd ("NFS handle COPY ERR_OFFLOAD_NO_REQS")
    Signed-off-by: Olga Kornievskaia <kolga@netapp.com>
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dd143ee28ef7dd62be46d15155e0ecac3c290d1a
Author: Benjamin Coddington <bcodding@redhat.com>
Date:   Tue Aug 22 14:22:38 2023 -0400

    NFS: Guard against READDIR loop when entry names exceed MAXNAMELEN

    [ Upstream commit f67b55b6588bcf9316a1e6e8d529100a5aa3ebe6 ]

    Commit 64cfca85bacd asserts the only valid return values for
    nfs2/3_decode_dirent should not include -ENAMETOOLONG, but for a server
    that sends a filename3 which exceeds MAXNAMELEN in a READDIR response the
    client's behavior will be to endlessly retry the operation.

    We could map -ENAMETOOLONG into -EBADCOOKIE, but that would produce
    truncated listings without any error.  The client should return an error
    for this case to clearly assert that the server implementation must be
    corrected.

    Fixes: 64cfca85bacd ("NFS: Return valid errors from nfs2/3_decode_dirent()")
    Signed-off-by: Benjamin Coddington <bcodding@redhat.com>
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 71ebbefeb9965adb5b1bd8da86e68bd8918df348
Author: Nathan Chancellor <nathan@kernel.org>
Date:   Tue Aug 29 07:08:47 2023 -0700

    clk: qcom: Fix SM_GPUCC_8450 dependencies

    [ Upstream commit 75d1d3a433f0a0748a89eb074830e9b635a19fd2 ]

    CONFIG_SM_GCC_8450 depends on ARM64 but it is selected by
    CONFIG_SM_GPUCC_8450, which can be selected on ARM, resulting in a
    Kconfig warning.

    WARNING: unmet direct dependencies detected for SM_GCC_8450
      Depends on [n]: COMMON_CLK [=y] && COMMON_CLK_QCOM [=y] && (ARM64 || COMPILE_TEST
[=n])
      Selected by [y]:
      - SM_GPUCC_8450 [=y] && COMMON_CLK [=y] && COMMON_CLK_QCOM [=y]

    Add the same dependencies to CONFIG_SM_GPUCC_8450 to resolve the
    warning.

    Fixes: 728692d49edc ("clk: qcom: Add support for SM8450 GPUCC")
    Signed-off-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20230829-fix-sm_gpucc_8550-deps-v1-1-
d751f6cd35b2@kernel.org
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64c959d59c24c661bd5338e4e6f1f703d7c309a1
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Wed Aug 16 10:20:52 2023 -0400

    NFSD: da_addr_body field missing in some GETDEVICEINFO replies

    [ Upstream commit 6372e2ee629894433fe6107d7048536a3280a284 ]

    The XDR specification in RFC 8881 looks like this:

    struct device_addr4 {
            layouttype4     da_layout_type;
            opaque          da_addr_body<>;
    };

    struct GETDEVICEINFO4resok {
            device_addr4    gdir_device_addr;
            bitmap4         gdir_notification;
    };

    union GETDEVICEINFO4res switch (nfsstat4 gdir_status) {
    case NFS4_OK:
            GETDEVICEINFO4resok gdir_resok4;
    case NFS4ERR_TOOSMALL:
            count4          gdir_mincount;
    default:
            void;
    };

    Looking at nfsd4_encode_getdeviceinfo() ....

    When the client provides a zero gd_maxcount, then the Linux NFS
    server implementation encodes the da_layout_type field and then
    skips the da_addr_body field completely, proceeding directly to
    encode gdir_notification field.
```

```
    There does not appear to be an option in the specification to skip
    encoding da_addr_body. Moreover, Section 18.40.3 says:

    > If the client wants to just update or turn off notifications, it
    > MAY send a GETDEVICEINFO operation with gdia_maxcount set to zero.
    > In that event, if the device ID is valid, the reply's da_addr_body
    > field of the gdir_device_addr field will be of zero length.

    Since the layout drivers are responsible for encoding the
    da_addr_body field, put this fix inside the ->encode_getdeviceinfo
    methods.

    Fixes: 9cf514ccfacb ("nfsd: implement pNFS operations")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Cc: Tom Haynes <loghyr@gmail.com>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 870a2dbc7bc94a7bac58d9a9bdfe308fae9619f3
Author: Su Hui <suhui@nfschina.com>
Date:   Fri Aug 4 09:26:57 2023 +0800

    fs: lockd: avoid possible wrong NULL parameter

    [ Upstream commit de8d38cf44bac43e83bad28357ba84784c412752 ]

    clang's static analysis warning: fs/lockd/mon.c: line 293, column 2:
    Null pointer passed as 2nd argument to memory copy function.

    Assuming 'hostname' is NULL and calling 'nsm_create_handle()', this will
    pass NULL as 2nd argument to memory copy function 'memcpy()'. So return
    NULL if 'hostname' is invalid.

    Fixes: 77a3ef33e2de ("NSM: More clean up of nsm_get_handle()")
    Signed-off-by: Su Hui <suhui@nfschina.com>
    Reviewed-by: Nick Desaulniers <ndesaulniers@google.com>
    Reviewed-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3bf8d3ca5a9299a4129fee40df6df22cf93d4c19
Author: Chuck Lever <chuck.lever@oracle.com>
Date:   Mon Aug 28 09:23:00 2023 -0400

    SUNRPC: Fix the recent bv_offset fix

    [ Upstream commit f16ff1cafbf1e65cc706af912df90bcc15d39a6c ]

    Jeff confirmed his original fix addressed his pynfs test failure,
    but this same bug also impacted qemu: accessing qcow2 virtual disks
    using direct I/O was failing. Jeff's fix missed that you have to
    shorten the bio_vec element by the same amount as you increased
    the page offset.

    Reported-by: Maxim Levitsky <mlevitsk@redhat.com>
    Fixes: c96e2a695e00 ("sunrpc: set the bv_offset of first bvec in svc_tcp_sendmsg")
    Tested-by: Maxim Levitsky <mlevitsk@redhat.com>
    Reviewed-by: Jeff Layton <jlayton@kernel.org>
    Signed-off-by: Chuck Lever <chuck.lever@oracle.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f7121a903aa7c8536749a6e4cdeb3044981a3ee2
Author: Alexei Filippov <halip0503@gmail.com>
Date:   Sat Aug 19 20:32:16 2023 +0300
```

```
    jfs: validate max amount of blocks before allocation.

    [ Upstream commit 0225e10972fa809728b8d4c1bd2772b3ec3fdb57 ]

    The lack of checking bmp->db_max_freebud in extBalloc() can lead to
    shift out of bounds, so this patch prevents undefined behavior, because
    bmp->db_max_freebud == -1 only if there is no free space.

    Signed-off-by: Aleksei Filippov <halip0503@gmail.com>
    Signed-off-by: Dave Kleikamp <dave.kleikamp@oracle.com>
    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reported-and-tested-by: syzbot+5f088f29593e6b4c8db8@syzkaller.appspotmail.com
    Closes: https://syzkaller.appspot.com/bug?id=01abadbd6ae6a08b1f1987aa61554c6b3ac19ff2
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8bfada8501b21d268aea5177ab0294d354bb03e
Author: Zhihao Cheng <chengzhihao1@huawei.com>
Date:   Wed Jun 28 21:20:11 2023 +0800

    ext4: fix unttached inode after power cut with orphan file feature enabled

    [ Upstream commit 1524773425ae8113b0b782886366e68656b34e53 ]

    Running generic/475(filesystem consistent tests after power cut) could
    easily trigger unattached inode error while doing fsck:
      Unattached zero-length inode 39405.  Clear? no

      Unattached inode 39405
      Connect to /lost+found? no

    Above inconsistence is caused by following process:
           P1                        P2
    ext4_create
     inode = ext4_new_inode_start_handle  // itable records nlink=1
     ext4_add_nondir
       err = ext4_add_entry  // ENOSPC
        ext4_append
         ext4_bread
          ext4_getblk
           ext4_map_blocks // returns ENOSPC
       drop_nlink(inode) // won't be updated into disk inode
       ext4_orphan_add(handle, inode)
        ext4_orphan_file_add
     ext4_journal_stop(handle)
                           jbd2_journal_commit_transaction // commit success
                  >> power cut <<
    ext4_fill_super
     ext4_load_and_init_journal   // itable records nlink=1
     ext4_orphan_cleanup
      ext4_process_orphan
       if (inode->i_nlink)        // true, inode won't be deleted

    Then, allocated inode will be reserved on disk and corresponds to no
    dentries, so e2fsck reports 'unattached inode' problem.

    The problem won't happen if orphan file feature is disabled, because
    ext4_orphan_add() will update disk inode in orphan list mode. There
    are several places not updating disk inode while putting inode into
    orphan area, such as ext4_add_nondir(), ext4_symlink() and whiteout
    in ext4_rename(). Fix it by updating inode into disk in all error
    branches of these places.

    Link: https://bugzilla.kernel.org/show_bug.cgi?id=217605
    Fixes: 02f310fcf47f ("ext4: Speedup ext4 orphan inode handling")
    Signed-off-by: Zhihao Cheng <chengzhihao1@huawei.com>
```

```
    Reviewed-by: Jan Kara <jack@suse.cz>
    Link: https://lore.kernel.org/r/20230628132011.650383-1-chengzhihao1@huawei.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6670c65bf863cd0d44ca24d4c10ef6755b8d9529
Author: Russell Currey <ruscur@russell.cc>
Date:   Wed Mar 22 14:53:22 2023 +1100

    powerpc/iommu: Fix notifiers being shared by PCI and VIO buses

    [ Upstream commit c37b6908f7b2bd24dcaaf14a180e28c9132b9c58 ]

    fail_iommu_setup() registers the fail_iommu_bus_notifier struct to both
    PCI and VIO buses.  struct notifier_block is a linked list node, so this
    causes any notifiers later registered to either bus type to also be
    registered to the other since they share the same node.

    This causes issues in (at least) the vgaarb code, which registers a
    notifier for PCI buses.  pci_notify() ends up being called on a vio
    device, converted with to_pci_dev() even though it's not a PCI device,
    and finally makes a bad access in vga_arbiter_add_pci_device() as
    discovered with KASAN:

     BUG: KASAN: slab-out-of-bounds in vga_arbiter_add_pci_device+0x60/0xe00
     Read of size 4 at addr c000000264c26fdc by task swapper/0/1

     Call Trace:
       dump_stack_lvl+0x1bc/0x2b8 (unreliable)
       print_report+0x3f4/0xc60
       kasan_report+0x244/0x698
       __asan_load4+0xe8/0x250
       vga_arbiter_add_pci_device+0x60/0xe00
       pci_notify+0x88/0x444
       notifier_call_chain+0x104/0x320
       blocking_notifier_call_chain+0xa0/0x140
       device_add+0xac8/0x1d30
       device_register+0x58/0x80
       vio_register_device_node+0x9ac/0xce0
       vio_bus_scan_register_devices+0xc4/0x13c
       __machine_initcall_pseries_vio_device_init+0x94/0xf0
       do_one_initcall+0x12c/0xaa8
       kernel_init_freeable+0xa48/0xba8
       kernel_init+0x64/0x400
       ret_from_kernel_thread+0x5c/0x64

    Fix this by creating separate notifier_block structs for each bus type.

    Fixes: d6b9a81b2a45 ("powerpc: IOMMU fault injection")
    Reported-by: Nageswara R Sastry <rnsastry@linux.ibm.com>
    Signed-off-by: Russell Currey <ruscur@russell.cc>
    Tested-by: Nageswara R Sastry <rnsastry@linux.ibm.com>
    Reviewed-by: Andrew Donnellan <ajd@linux.ibm.com>
    [mpe: Add #ifdef to fix CONFIG_IBMVIO=n build]
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230322035322.328709-1-ruscur@russell.cc
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b623e1a4bd05671d41750f4c439ea5786f2abe4
Author: Liang He <windhl@126.com>
Date:   Wed Mar 22 11:04:23 2023 +0800

    powerpc/mpc5xxx: Add missing fwnode_handle_put()

    [ Upstream commit b9bbbf4979073d5536b7650decd37fcb901e6556 ]
```

```
    In mpc5xxx_fwnode_get_bus_frequency(), we should add
    fwnode_handle_put() when break out of the iteration
    fwnode_for_each_parent_node() as it will automatically
    increase and decrease the refcounter.

    Fixes: de06fba62af6 ("powerpc/mpc5xxx: Switch mpc5xxx_get_bus_frequency() to use
fwnode")
    Signed-off-by: Liang He <windhl@126.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230322030423.1855440-1-windhl@126.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac6d060d97092d766485f62f5b2b4987b73f79a7
Author: Nicholas Piggin <npiggin@gmail.com>
Date:   Tue May 9 19:15:59 2023 +1000

    powerpc/pseries: Fix hcall tracepoints with JUMP_LABEL=n

    [ Upstream commit 750bd41aeaeb1f0e0128aa4f8fcd6dd759713641 ]

    With JUMP_LABEL=n, hcall_tracepoint_refcount's address is being tested
    instead of its value. This results in the tracing slowpath always being
    taken unnecessarily.

    Fixes: 9a10ccb29c0a2 ("powerpc/pseries: move hcall_tracepoint_refcount out of .toc")
    Signed-off-by: Nicholas Piggin <npiggin@gmail.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230509091600.70994-1-npiggin@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d480874e8264b39e609212c99f534040713b70a
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Mon Jul 24 11:08:46 2023 +0300

    nfs/blocklayout: Use the passed in gfp flags

    [ Upstream commit 08b45fcb2d4675f6182fe0edc0d8b1fe604051fa ]

    This allocation should use the passed in GFP_ flags instead of
    GFP_KERNEL.  One places where this matters is in filelayout_pg_init_write()
    which uses GFP_NOFS as the allocation flags.

    Fixes: 5c83746a0cf2 ("pnfs/blocklayout: in-kernel GETDEVICEINFO XDR parsing")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f45ee5c074013a0fbfce77a5af5efddb01f5d4f4
Author: Russell Currey <ruscur@russell.cc>
Date:   Wed Aug 23 15:53:17 2023 +1000

    powerpc/pseries: Rework lppaca_shared_proc() to avoid DEBUG_PREEMPT

    [ Upstream commit eac030b22ea12cdfcbb2e941c21c03964403c63f ]

    lppaca_shared_proc() takes a pointer to the lppaca which is typically
    accessed through get_lppaca().  With DEBUG_PREEMPT enabled, this leads
    to checking if preemption is enabled, for example:

      BUG: using smp_processor_id() in preemptible [00000000] code: grep/10693
      caller is lparcfg_data+0x408/0x19a0
      CPU: 4 PID: 10693 Comm: grep Not tainted 6.5.0-rc3 #2
      Call Trace:
```

```
         dump_stack_lvl+0x154/0x200 (unreliable)
         check_preemption_disabled+0x214/0x220
         lparcfg_data+0x408/0x19a0
         ...

    This isn't actually a problem however, as it does not matter which
    lppaca is accessed, the shared proc state will be the same.
    vcpudispatch_stats_procfs_init() already works around this by disabling
    preemption, but the lparcfg code does not, erroring any time
    /proc/powerpc/lparcfg is accessed with DEBUG_PREEMPT enabled.

    Instead of disabling preemption on the caller side, rework
    lppaca_shared_proc() to not take a pointer and instead directly access
    the lppaca, bypassing any potential preemption checks.

    Fixes: f13c13a00512 ("powerpc: Stop using non-architected shared_proc field in lppaca")
    Signed-off-by: Russell Currey <ruscur@russell.cc>
    [mpe: Rework to avoid needing a definition in paca.h and lppaca.h]
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230823055317.751786-4-mpe@ellerman.id.au
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89d8cfd1256648dd554f6ad3fa943b71d7478cbe
Author: Michael Ellerman <mpe@ellerman.id.au>
Date:   Wed Aug 23 15:53:16 2023 +1000

    powerpc: Don't include lppaca.h in paca.h

    [ Upstream commit 1aa000667669fa855853decbb1c69e974d8ff716 ]

    By adding a forward declaration for struct lppaca we can untangle paca.h
    and lppaca.h. Also move get_lppaca() into lppaca.h for consistency.

    Add includes of lppaca.h to some files that need it.

    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230823055317.751786-3-mpe@ellerman.id.au
    Stable-dep-of: eac030b22ea1 ("powerpc/pseries: Rework lppaca_shared_proc() to avoid
DEBUG_PREEMPT")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 517868d1a334ab4b8d7af3a60642924814317658
Author: Xiaowei Bao <xiaowei.bao@nxp.com>
Date:   Thu Jul 20 09:58:34 2023 -0400

    PCI: layerscape: Add workaround for lost link capabilities during reset

    [ Upstream commit 17cf8661ee0f065c08152e611a568dd1fb0285f1 ]

    The endpoint controller loses the Maximum Link Width and Supported Link Speed
    value from the Link Capabilities Register - initially configured by the Reset
    Configuration Word (RCW) - during a link-down or hot reset event.

    Address this issue in the endpoint event handler.

    Link: https://lore.kernel.org/r/20230720135834.1977616-2-Frank.Li@nxp.com
    Fixes: a805770d8a22 ("PCI: layerscape: Add EP mode support")
    Signed-off-by: Xiaowei Bao <xiaowei.bao@nxp.com>
    Signed-off-by: Hou Zhiqiang <Zhiqiang.Hou@nxp.com>
    Signed-off-by: Frank Li <Frank.Li@nxp.com>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Acked-by: Manivannan Sadhasivam <mani@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae5d5672f1db711e91db6f52df5cb16ecd8f5692
```

```
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Fri Jun 9 15:26:25 2023 -0400

    NFSv4.2: Rework scratch handling for READ_PLUS (again)

    [ Upstream commit 303a78052091c81e9003915c521fdca1c7e117af ]

    I found that the read code might send multiple requests using the same
    nfs_pgio_header, but nfs4_proc_read_setup() is only called once. This is
    how we ended up occasionally double-freeing the scratch buffer, but also
    means we set a NULL pointer but non-zero length to the xdr scratch
    buffer. This results in an oops the first time decoding needs to copy
    something to scratch, which frequently happens when decoding READ_PLUS
    hole segments.

    I fix this by moving scratch handling into the pageio read code. I
    provide a function to allocate scratch space for decoding read replies,
    and free the scratch buffer when the nfs_pgio_header is freed.

    Fixes: fbd2a05f29a9 (NFSv4.2: Rework scratch handling for READ_PLUS)
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a15c9dcfe766232f731f98fecce133e1d38f73ea
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Wed May 31 17:02:54 2023 -0400

    NFSv4.2: Fix READ_PLUS size calculations

    [ Upstream commit 8d18f6c5bb864d97a730f471c56cdecf313efe64 ]

    I bump the decode_read_plus_maxsz to account for hole segments, but I
    need to subtract out this increase when calling
    rpc_prepare_reply_pages() so the common case of single data segment
    replies can be directly placed into the xdr pages without needing to be
    shifted around.

    Reported-by: Chuck Lever <chuck.lever@oracle.com>
    Fixes: d3b00a802c845 ("NFS: Replace the READ_PLUS decoding code")
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5479d25eaeeab9d13669971bbfe9606329581692
Author: Anna Schumaker <Anna.Schumaker@Netapp.com>
Date:   Wed May 24 17:27:08 2023 -0400

    NFSv4.2: Fix READ_PLUS smatch warnings

    [ Upstream commit bb05a617f06b7a882e19c4f475b8e37f14d9ceac ]

    Smatch reports:
      fs/nfs/nfs42xdr.c:1131 decode_read_plus() warn: missing error code? 'status'

    Which Dan suggests to fix by doing a hardcoded "return 0" from the
    "if (segments == 0)" check.

    Additionally, smatch reports that the "status = -EIO" assignment is not
    used. This patch addresses both these issues.

    Reported-by: kernel test robot <lkp@intel.com>
    Reported-by: Dan Carpenter <error27@gmail.com>
    Closes: https://lore.kernel.org/r/202305222209.6l5VM2lL-lkp@intel.com/
    Fixes: d3b00a802c845 ("NFS: Replace the READ_PLUS decoding code")
    Signed-off-by: Anna Schumaker <Anna.Schumaker@Netapp.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 8c2314bd89c74bdc9b454691e11643a00b575a72
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:05:02 2023 +0300

    wifi: ath10k: Use RMW accessors for changing LNKCTL

    [ Upstream commit f139492a09f15254fa261245cdbd65555cdf39e3 ]

    Don't assume that only the driver would be accessing LNKCTL. ASPM policy
    changes can trigger write to LNKCTL outside of driver's control.

    Use RMW capability accessors which does proper locking to avoid losing
    concurrent updates to the register value. On restore, clear the ASPMC field
    properly.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: 76d870ed09ab ("ath10k: enable ASPM")
    Link: https://lore.kernel.org/r/20230717120503.15276-11-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Acked-by: Kalle Valo <kvalo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0639558df9fd1c27c578d8cf4cc1dc8f73697cfb
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:05:01 2023 +0300

    wifi: ath12k: Use RMW accessors for changing LNKCTL

    [ Upstream commit f5a7ac118faf6d4f794975947b3300717eae8fc5 ]

    Don't assume that only the driver would be accessing LNKCTL. ASPM policy
    changes can trigger write to LNKCTL outside of driver's control.

    Use RMW capability accessors which do proper locking to avoid losing
    concurrent updates to the register value. On restore, clear the ASPMC field
    properly.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: d889913205cf ("wifi: ath12k: driver for Qualcomm Wi-Fi 7 devices")
    Link: https://lore.kernel.org/r/20230717120503.15276-10-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Acked-by: Kalle Valo <kvalo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 981ab31010653e86a91d14b27fa900257dbe9bc5
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:05:00 2023 +0300

    wifi: ath11k: Use RMW accessors for changing LNKCTL

    [ Upstream commit 6c1b6bdb34aaf8f94f65a9cae1d63490320c11bc ]

    Don't assume that only the driver would be accessing LNKCTL. ASPM policy
    changes can trigger write to LNKCTL outside of driver's control.

    Use RMW capability accessors which do proper locking to avoid losing
    concurrent updates to the register value. On restore, clear the ASPMC field
    properly.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
```

```
    Fixes: e9603f4bdcc0 ("ath11k: pci: disable ASPM L0sLs before downloading firmware")
    Link: https://lore.kernel.org/r/20230717120503.15276-9-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Acked-by: Kalle Valo <kvalo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0653830e9356843cf161278d8dc687ec1b33dc87
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:59 2023 +0300

    net/mlx5: Use RMW accessors for changing LNKCTL

    [ Upstream commit 30de872537bda526664d7a20b646adfb3e7ce6e6 ]

    Don't assume that only the driver would be accessing LNKCTL of the upstream
    bridge. ASPM policy changes can trigger write to LNKCTL outside of driver's
    control.

    Use RMW capability accessors which do proper locking to avoid losing
    concurrent updates to the register value.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: eabe8e5e88f5 ("net/mlx5: Handle sync reset now event")
    Link: https://lore.kernel.org/r/20230717120503.15276-8-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Moshe Shemesh <moshe@nvidia.com>
    Reviewed-by: Simon Horman <simon.horman@corigine.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8976f9fd3caeb6ed29abf45fbe43974ffa0e7286
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:58 2023 +0300

    drm/radeon: Use RMW accessors for changing LNKCTL

    [ Upstream commit 7189576e8a829130192b33c5b64e8a475369c776 ]

    Don't assume that only the driver would be accessing LNKCTL. ASPM policy
    changes can trigger write to LNKCTL outside of driver's control.  And in
    the case of upstream bridge, the driver does not even own the device it's
    changing the registers for.

    Use RMW capability accessors which do proper locking to avoid losing
    concurrent updates to the register value.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: 8a7cd27679d0 ("drm/radeon/cik: add support for pcie gen1/2/3 switching")
    Fixes: b9d305dfb66c ("drm/radeon: implement pcie gen2/3 support for SI")
    Link: https://lore.kernel.org/r/20230717120503.15276-7-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8aaed0f373b9f7fa80e6c92ee2b8677f81739f79
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:57 2023 +0300

    drm/amdgpu: Use RMW accessors for changing LNKCTL

    [ Upstream commit ce7d88110b9ed5f33fe79ea6d4ed049fb0e57bce ]
```

```
    Don't assume that only the driver would be accessing LNKCTL. ASPM policy
    changes can trigger write to LNKCTL outside of driver's control.  And in
    the case of upstream bridge, the driver does not even own the device it's
    changing the registers for.

    Use RMW capability accessors which do proper locking to avoid losing
    concurrent updates to the register value.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: a2e73f56fa62 ("drm/amdgpu: Add support for CIK parts")
    Fixes: 62a37553414a ("drm/amdgpu: add si implementation v10")
    Link: https://lore.kernel.org/r/20230717120503.15276-6-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ffadca11bd814030f47e6a8995c8469bacbc8e3c
Author: Daniel Golle <daniel@makrotopia.org>
Date:   Fri Aug 18 04:03:26 2023 +0100

    pinctrl: mediatek: assign functions to configure pin bias on MT7986

    [ Upstream commit 0d8387fba9f151220e48dc3dcdc2335539708f13 ]

    Assign bias_disable_get/set and bias_get/set functions to allow
    configuring pin bias on MT7986.

    Fixes: 2c58d8dc9cd0 ("pinctrl: mediatek: add pull_type attribute for mediatek MT7986
SoC")
    Signed-off-by: Daniel Golle <daniel@makrotopia.org>
    Link:
https://lore.kernel.org/r/47f72372354312a839b9337e09476aadcc206e8b.1692327317.git.daniel@mak
rotopia.org
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a601e9a8ada59d33872c9fb2ce58dc8dca750f53
Author: Daniel Golle <daniel@makrotopia.org>
Date:   Fri Aug 18 04:02:35 2023 +0100

    pinctrl: mediatek: fix pull_type data for MT7981

    [ Upstream commit 8f6f16fe1553ce63edfb98a39ef9d4754a0c39bf ]

    MediaTek has released pull_type data for MT7981 in their SDK.
    Use it and set functions to configure pin bias.

    Fixes: 6c83b2d94fcc ("pinctrl: add mt7981 pinctrl driver")
    Signed-off-by: Daniel Golle <daniel@makrotopia.org>
    Link:
https://lore.kernel.org/r/7bcc8ead25dbfabc7f5a85d066224a926fbb4941.1692327317.git.daniel@mak
rotopia.org
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3aef48afc93780ccc5fe3d604a87b68b06f0d2d
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Fri Aug 18 10:59:44 2023 +0200

    powerpc/perf: Convert fsl_emb notifier to state machine callbacks

    [ Upstream commit 34daf445f82bd3a4df852bb5f1dffd792ac830a0 ]

      CC      arch/powerpc/perf/core-fsl-emb.o
```

```
    arch/powerpc/perf/core-fsl-emb.c:675:6: error: no previous prototype for
 'hw_perf_event_setup' [-Werror=missing-prototypes]
      675 | void hw_perf_event_setup(int cpu)
          |      ^~~~~~~~~~~~~~~~~~

    Looks like fsl_emb was completely missed by commit 3f6da3905398 ("perf:
    Rework and fix the arch CPU-hotplug hooks")

    So, apply same changes as commit 3f6da3905398 ("perf: Rework and fix
    the arch CPU-hotplug hooks") then commit 57ecde42cc74 ("powerpc/perf:
    Convert book3s notifier to state machine callbacks")

    While at it, also fix following error:

    arch/powerpc/perf/core-fsl-emb.c: In function 'perf_event_interrupt':
    arch/powerpc/perf/core-fsl-emb.c:648:13: error: variable 'found' set but not used [-
 Werror=unused-but-set-variable]
      648 |         int found = 0;
          |             ^~~~

    Fixes: 3f6da3905398 ("perf: Rework and fix the arch CPU-hotplug hooks")
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link:
 https://msgid.link/603e1facb32608f88f40b7d7b9094adc50e7b2dc.1692349125.git.christophe.leroy@
 csgroup.eu
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 7cad534400887effa1d90d165c1b079617e9cc80
 Author: Sourabh Jain <sourabhjain@linux.ibm.com>
 Date:   Tue Jul 4 10:37:15 2023 +0530

    powerpc/fadump: reset dump area size if fadump memory reserve fails

    [ Upstream commit d1eb75e0dfed80d2d85b664e28a39f65b290ab55 ]

    In case fadump_reserve_mem() fails to reserve memory, the
    reserve_dump_area_size variable will retain the reserve area size. This
    will lead to /sys/kernel/fadump/mem_reserved node displaying an incorrect
    memory reserved by fadump.

    To fix this problem, reserve dump area size variable is set to 0 if fadump
    failed to reserve memory.

    Fixes: 8255da95e545 ("powerpc/fadump: release all the memory above boot memory size")
    Signed-off-by: Sourabh Jain <sourabhjain@linux.ibm.com>
    Acked-by: Mahesh Salgaonkar <mahesh@linux.ibm.com>
    Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
    Link: https://msgid.link/20230704050715.203581-1-sourabhjain@linux.ibm.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit cebe779e5eee84fb375c8163d720387392cdf9a3
 Author: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
 Date:   Thu Aug 17 19:41:03 2023 +0800

    nvdimm: Fix dereference after free in register_nvdimm_pmu()

    [ Upstream commit 08ca6906a4b7e48f8e93b7c1f49a742a415be6d5 ]

    'nd_pmu->pmu.attr_groups' is dereferenced in function
    'nvdimm_pmu_free_hotplug_memory' call after it has been freed. Because in
    function 'nvdimm_pmu_free_hotplug_memory' memory pointed by the fields of
    'nd_pmu->pmu.attr_groups' is deallocated it is necessary to call 'kfree'
    after 'nvdimm_pmu_free_hotplug_memory'.
```

```
    Fixes: 0fab1ba6ad6b ("drivers/nvdimm: Add perf interface to expose nvdimm performance
stats")
    Co-developed-by: Ivanov Mikhail <ivanov.mikhail1@huawei-partners.com>
    Signed-off-by: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
    Reviewed-by: Jeff Moyer <jmoyer@redhat.com>
    Link: https://lore.kernel.org/r/20230817114103.754977-1-konstantin.meskhidze@huawei.com
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 16259c80542ee8945aaa39cfc6a1809bcdc08ffe
Author: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
Date:   Thu Aug 17 19:59:45 2023 +0800

    nvdimm: Fix memleak of pmu attr_groups in unregister_nvdimm_pmu()

    [ Upstream commit 85ae42c72142346645e63c33835da947dfa008b3 ]

    Memory pointed by 'nd_pmu->pmu.attr_groups' is allocated in function
    'register_nvdimm_pmu' and is lost after 'kfree(nd_pmu)' call in function
    'unregister_nvdimm_pmu'.

    Fixes: 0fab1ba6ad6b ("drivers/nvdimm: Add perf interface to expose nvdimm performance
stats")
    Co-developed-by: Ivanov Mikhail <ivanov.mikhail1@huawei-partners.com>
    Signed-off-by: Konstantin Meskhidze <konstantin.meskhidze@huawei.com>
    Reviewed-by: Jeff Moyer <jmoyer@redhat.com>
    Link: https://lore.kernel.org/r/20230817115945.771826-1-konstantin.meskhidze@huawei.com
    Signed-off-by: Dave Jiang <dave.jiang@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b5feb8497cdb5b9962db2700814bffbc030fb4a
Author: Stefan Hajnoczi <stefanha@redhat.com>
Date:   Tue Aug 1 11:53:52 2023 -0400

    vfio/type1: fix cap_migration information leak

    [ Upstream commit cd24e2a60af633f157d7e59c0a6dba64f131c0b1 ]

    Fix an information leak where an uninitialized hole in struct
    vfio_iommu_type1_info_cap_migration on the stack is exposed to userspace.

    The definition of struct vfio_iommu_type1_info_cap_migration contains a hole as
    shown in this pahole(1) output:

      struct vfio_iommu_type1_info_cap_migration {
              struct vfio_info_cap_header header;              /*     0     8 */
              __u32                       flags;               /*     8     4 */

              /* XXX 4 bytes hole, try to pack */

              __u64                       pgsize_bitmap;       /*    16     8 */
              __u64                       max_dirty_bitmap_size; /*    24     8 */

              /* size: 32, cachelines: 1, members: 4 */
              /* sum members: 28, holes: 1, sum holes: 4 */
              /* last cacheline: 32 bytes */
      };

    The cap_mig variable is filled in without initializing the hole:

      static int vfio_iommu_migration_build_caps(struct vfio_iommu *iommu,
                              struct vfio_info_cap *caps)
      {
          struct vfio_iommu_type1_info_cap_migration cap_mig;
```

```
            cap_mig.header.id = VFIO_IOMMU_TYPE1_INFO_CAP_MIGRATION;
            cap_mig.header.version = 1;

            cap_mig.flags = 0;
            /* support minimum pgsize */
            cap_mig.pgsize_bitmap = (size_t)1 << __ffs(iommu->pgsize_bitmap);
            cap_mig.max_dirty_bitmap_size = DIRTY_BITMAP_SIZE_MAX;

            return vfio_info_add_capability(caps, &cap_mig.header, sizeof(cap_mig));
        }
```

    The structure is then copied to a temporary location on the heap. At this point
    it's already too late and ioctl(VFIO_IOMMU_GET_INFO) copies it to userspace
    later:

```
        int vfio_info_add_capability(struct vfio_info_cap *caps,
                        struct vfio_info_cap_header *cap, size_t size)
        {
            struct vfio_info_cap_header *header;

            header = vfio_info_cap_add(caps, size, cap->id, cap->version);
            if (IS_ERR(header))
                return PTR_ERR(header);

            memcpy(header + 1, cap + 1, size - sizeof(*header));

            return 0;
        }
```

    This issue was found by code inspection.

    Signed-off-by: Stefan Hajnoczi <stefanha@redhat.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Fixes: ad721705d09c ("vfio iommu: Add migration capability to report supported
features")
    Link: https://lore.kernel.org/r/20230801155352.1391945-1-stefanha@redhat.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e62c6b05dda8ff126d2db354bdeefc586128576d
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date: Wed Aug 9 10:01:43 2023 +0200

    powerpc/radix: Move some functions into #ifdef CONFIG_KVM_BOOK3S_HV_POSSIBLE

    [Upstream commit 4a9dd8f292efd614f0a18452e6474fe19ae17b47]

    With skiboot_defconfig, Clang reports:

```
      CC      arch/powerpc/mm/book3s64/radix_tlb.o
    arch/powerpc/mm/book3s64/radix_tlb.c:419:20: error: unused function '_tlbie_pid_lpid' [-
Werror,-Wunused-function]
    static inline void _tlbie_pid_lpid(unsigned long pid, unsigned long lpid,
                       ^
    arch/powerpc/mm/book3s64/radix_tlb.c:663:20: error: unused function
'_tlbie_va_range_lpid' [-Werror,-Wunused-function]
    static inline void _tlbie_va_range_lpid(unsigned long start, unsigned long end,
                       ^
```

    This is because those functions are only called from functions
    enclosed in a #ifdef CONFIG_KVM_BOOK3S_HV_POSSIBLE

    Move below functions inside that #ifdef
    * __tlbie_pid_lpid(unsigned long pid,
    * __tlbie_va_lpid(unsigned long va, unsigned long pid,

```
      * fixup_tlbie_pid_lpid(unsigned long pid, unsigned long lpid)
      * _tlbie_pid_lpid(unsigned long pid, unsigned long lpid,
      * fixup_tlbie_va_range_lpid(unsigned long va,
      * __tlbie_va_range_lpid(unsigned long start, unsigned long end,
      * _tlbie_va_range_lpid(unsigned long start, unsigned long end,

      Fixes: f0c6fbbb9050 ("KVM: PPC: Book3S HV: Add support for H_RPT_INVALIDATE")
      Reported-by: kernel test robot <lkp@intel.com>
      Closes: https://lore.kernel.org/oe-kbuild-all/202307260802.Mjr99P5O-lkp@intel.com/
      Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
      Signed-off-by: Michael Ellerman <mpe@ellerman.id.au>
      Link:
  https://msgid.link/3d72efd39f986ee939d068af69fdce28bd600766.1691568093.git.christophe.leroy@
  csgroup.eu
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit fba6211b1fdf7c8655444c5cc608a1dcc7373022
  Author: Ahmad Fatoum <a.fatoum@pengutronix.de>
  Date:   Mon Aug 7 10:22:00 2023 +0200

      clk: imx: composite-8m: fix clock pauses when set_rate would be a no-op

      [ Upstream commit 4dd432d985ef258e3bc436e568fba4b987b59171 ]

      Reconfiguring the clock divider to the exact same value is observed
      on an i.MX8MN to often cause a longer than usual clock pause, probably
      because the divider restarts counting whenever the register is rewritten.

      This issue doesn't show up normally, because the clock framework will
      take care to not call set_rate when the clock rate is the same.
      However, when we reconfigure an upstream clock, the common code will
      call set_rate with the newly calculated rate on all children, e.g.:

        - sai5 is running normally and divides Audio PLL out by 16.
        - Audio PLL rate is increased by 32Hz (glitch-free kdiv change)
        - rates for children are recalculated and rates are set recursively
        - imx8m_clk_composite_divider_set_rate(sai5) is called with
          32/16 = 2Hz more
        - imx8m_clk_composite_divider_set_rate computes same divider as before
        - divider register is written, so it restarts counting from zero and
          MCLK is briefly paused, so instead of e.g. 40ns, MCLK is low for 120ns.

      Some external clock consumers can be upset by such unexpected clock pauses,
      so let's make sure we only rewrite the divider value when the value to be
      written is actually different.

      Fixes: d3ff9728134e ("clk: imx: Add imx composite clock")
      Signed-off-by: Ahmad Fatoum <a.fatoum@pengutronix.de>
      Reviewed-by: Peng Fan <peng.fan@nxp.com>
      Link: https://lore.kernel.org/r/20230807082201.2332746-1-a.fatoum@pengutronix.de
      Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 04295540d1f81c0b2db2474a22e24dde768ee13e
  Author: Marco Felsch <m.felsch@pengutronix.de>
  Date:   Mon Jul 31 16:21:49 2023 +0200

      clk: imx8mp: fix sai4 clock

      [ Upstream commit c30f600f1f41dcf5ef0fb02e9a201f9b2e8f31bd ]

      The reference manual don't mention a SAI4 hardware block. This would be
      clock slice 78 which is skipped (TRM, page 237). Remove any reference to
      this clock to align the driver with the reality.
```

```
    Fixes: 9c140d992676 ("clk: imx: Add support for i.MX8MP clock driver")
    Acked-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Link: https://lore.kernel.org/r/20230731142150.3186650-1-m.felsch@pengutronix.de
    Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 913c74d6ce0af589129beb442ddd94ceee5c995b
Author: Peng Fan <peng.fan@nxp.com>
Date:   Sun Jun 25 20:33:40 2023 +0800

    clk: imx: imx8ulp: update SPLL2 type

    [ Upstream commit 7653a59be8af043adc4c09473975a860e6055ff9 ]

    The SPLL2 on iMX8ULP is different with other frac PLLs, it can
    support VCO from 650Mhz to 1Ghz. Following the changes to pllv4,
    use the new type IMX_PLLV4_IMX8ULP_1GHZ.

    Fixes: c43a801a5789 ("clk: imx: Add clock driver for imx8ulp")
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Abel Vesa <abel.vesa@linaro.org>
    Link: https://lore.kernel.org/r/20230625123340.4067536-2-peng.fan@oss.nxp.com
    Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 980ff1dd429e0960c88b3377af6d28cf7ca229dc
Author: Ye Li <ye.li@nxp.com>
Date:   Sun Jun 25 20:33:39 2023 +0800

    clk: imx: pllv4: Fix SPLL2 MULT range

    [ Upstream commit 3f0cdb945471f1abd1cf4d172190e9c489c5052a ]

    The SPLL2 on iMX8ULP is different with other frac PLLs, it can
    support VCO from 650Mhz to 1Ghz. According to RM, the MULT is
    using a range from 27 to 54, not some fixed values. If using
    current PLL implementation, some clock rate can't be supported.

    Fix the issue by adding new type for the SPLL2 and use MULT range
    to replace MULT table

    Fixes: 5f0601c47c33 ("clk: imx: Update the pllv4 to support imx8ulp")
    Reviewed-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Jacky Bai <ping.bai@nxp.com>
    Signed-off-by: Ye Li <ye.li@nxp.com>
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    Reviewed-by: Abel Vesa <abel.vesa@linaro.org>
    Link: https://lore.kernel.org/r/20230625123340.4067536-1-peng.fan@oss.nxp.com
    Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 05fe9752b5c33d8672f949305f87db06456bdb96
Author: Imran Shaik <quic_imrashai@quicinc.com>
Date:   Thu Aug 3 16:27:36 2023 +0530

    clk: qcom: gcc-qdu1000: Fix clkref clocks handling

    [ Upstream commit 2524dae5cd453ca39e8ba1b95c2755a8a2d94059 ]

    Update the GCC clkref clock's halt_check to BRANCH_HALT, as it's
    status bit is not inverted in the latest hardware version of QDU1000
    and QRU1000 SoCs. While at it, fix the gcc clkref clock ops as well.

    Fixes: 1c9efb0bc040 ("clk: qcom: Add QDU1000 and QRU1000 GCC support")
```

```
    Signed-off-by: Imran Shaik <quic_imrashai@quicinc.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230803105741.2292309-4-quic_imrashai@quicinc.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b26ea635811510d1f90212c06058629da5955f5a
Author: Imran Shaik <quic_imrashai@quicinc.com>
Date:   Thu Aug 3 16:27:35 2023 +0530

    clk: qcom: gcc-qdu1000: Fix gcc_pcie_0_pipe_clk_src clock handling

    [ Upstream commit b311f5d3c4749259043a9a458a8db07915210142 ]

    Fix the gcc pcie pipe clock handling as per the clk_regmap_phy_mux_ops
    implementation to let the clock framework automatically park the clock
    at XO when the clock is switched off and restore the parent when the
    clock is switched on.

    Fixes: 1c9efb0bc040 ("clk: qcom: Add QDU1000 and QRU1000 GCC support")
    Co-developed-by: Taniya Das <quic_tdas@quicinc.com>
    Signed-off-by: Taniya Das <quic_tdas@quicinc.com>
    Signed-off-by: Imran Shaik <quic_imrashai@quicinc.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230803105741.2292309-3-quic_imrashai@quicinc.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b50b60d6000fa446dac88dfb44f01e1ab84282a
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Aug 11 19:35:53 2023 +0200

    clk: qcom: gcc-sm8450: Use floor ops for SDCC RCGs

    [ Upstream commit a27ac3806b0a0e6954fb5967223b8635242e5b8f ]

    Use the floor ops to prevent warnings like this at suspend exit and boot:

    mmc0: Card appears overclocked; req 800000 Hz, actual 25000000 Hz

    Fixes: db0c944ee92b ("clk: qcom: Add clock driver for SM8450")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Link: https://lore.kernel.org/r/20230811-topic-8450_clk-v1-1-88031478d548@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3fd9b669154a89f19019797a495fe7af6b5c2fb0
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:56 2023 +0300

    PCI/ASPM: Use RMW accessors for changing LNKCTL

    [ Upstream commit e09060b3b6b4661278ff8e1b7b81a37d5ea86eae ]

    Don't assume that the device is fully under the control of ASPM and use RMW
    capability accessors which do proper locking to avoid losing concurrent
    updates to the register values.

    If configuration fails in pcie_aspm_configure_common_clock(), the
    function attempts to restore the old PCI_EXP_LNKCTL_CCC settings. Store
    only the old PCI_EXP_LNKCTL_CCC bit for the relevant devices rather
    than the content of the whole LNKCTL registers. It aligns better with
    how pcie_lnkctl_clear_and_set() expects its parameter and makes the
    code more obvious to understand.
```

```
    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: 2a42d9dba784 ("PCIe: ASPM: Break out of endless loop waiting for PCI config bits
to switch")
    Fixes: 7d715a6c1ae5 ("PCI: add PCI Express ASPM support")
    Link: https://lore.kernel.org/r/20230717120503.15276-5-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: "Rafael J. Wysocki" <rafael@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0aa88a8bfe318f844e220ee2cdbd3d89596a33b4
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:55 2023 +0300

    PCI: pciehp: Use RMW accessors for changing LNKCTL

    [ Upstream commit 5f75f96c61039151c193775d776fde42477eace1 ]

    As hotplug is not the only driver touching LNKCTL, use the RMW capability
    accessor which handles concurrent changes correctly.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: 7f822999e12a ("PCI: pciehp: Add Disable/enable link functions")
    Link: https://lore.kernel.org/r/20230717120503.15276-4-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: "Rafael J. Wysocki" <rafael@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 54217659075b1d11cad3f8bfa08ebb64cb26cd49
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:54 2023 +0300

    PCI: Make link retraining use RMW accessors for changing LNKCTL

    [ Upstream commit fb0171a4c01b4825e36a5584eaa84291179c64ce ]

    Don't assume that the device is fully under the control of PCI core.  Use
    RMW capability accessors in link retraining which do proper locking to
    avoid losing concurrent updates to the register values.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: 4ec73791a64b ("PCI: Work around Pericom PCIe-to-PCI bridge Retrain Link erratum")
    Fixes: 7d715a6c1ae5 ("PCI: add PCI Express ASPM support")
    Link: https://lore.kernel.org/r/20230717120503.15276-3-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Acked-by: "Rafael J. Wysocki" <rafael@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47e7a6529b5e0626d76f45432c7710073e5b6a44
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 15:04:53 2023 +0300

    PCI: Add locking to RMW PCI Express Capability Register accessors

    [ Upstream commit 5e70d0acf0825f4390797360803503 71f8d6699a ]

    Many places in the kernel write the Link Control and Root Control PCI
    Express Capability Registers without proper concurrency control and this
    could result in losing the changes one of the writers intended to make.

    Add pcie_cap_lock spinlock into the struct pci_dev and use it to protect
    bit changes made in the RMW capability accessors. Protect only a selected
```

```
    set of registers by differentiating the RMW accessor internally to
    locked/unlocked variants using a wrapper which has the same signature as
    pcie_capability_clear_and_set_word(). As the Capability Register (pos)
    given to the wrapper is always a constant, the compiler should be able to
    simplify all the dead-code away.

    So far only the Link Control Register (ASPM, hotplug, link retraining,
    various drivers) and the Root Control Register (AER & PME) seem to
    require RMW locking.

    Suggested-by: Lukas Wunner <lukas@wunner.de>
    Fixes: c7f486567c1d ("PCI PM: PCIe PME root port service driver")
    Fixes: f12eb72a268b ("PCI/ASPM: Use PCI Express Capability accessors")
    Fixes: 7d715a6c1ae5 ("PCI: add PCI Express ASPM support")
    Fixes: affa48de8417 ("staging/rdma/hfi1: Add support for enabling/disabling PCIe ASPM")
    Fixes: 849a9366cba9 ("misc: rtsx: Add support new chip rts5228 mmc: rtsx: Add support
MMC_CAP2_NO_MMC")
    Fixes: 3d1e7aa80d1c ("misc: rtsx: Use pcie_capability_clear_and_set_word() for
PCI_EXP_LNKCTL")
    Fixes: c0e5f4e73a71 ("misc: rtsx: Add support for RTS5261")
    Fixes: 3df4fce739e2 ("misc: rtsx: separate aspm mode into MODE_REG and MODE_CFG")
    Fixes: 121e9c6b5c4c ("misc: rtsx: modify and fix init_hw function")
    Fixes: 19f3bd548f27 ("mfd: rtsx: Remove LCTLR defination")
    Fixes: 773ccdfd9cc6 ("mfd: rtsx: Read vendor setting from config space")
    Fixes: 8275b77a1513 ("mfd: rts5249: Add support for RTS5250S power saving")
    Fixes: 5da4e04ae480 ("misc: rtsx: Add support for RTS5260")
    Fixes: 0f49bfbd0f2e ("tg3: Use PCI Express Capability accessors")
    Fixes: 5e7dfd0fb94a ("tg3: Prevent corruption at 10 / 100Mbps w CLKREQ")
    Fixes: b726e493e8dc ("r8169: sync existing 8168 device hardware start sequences with
vendor driver")
    Fixes: e6de30d63eb1 ("r8169: more 8168dp support.")
    Fixes: 8a06127602de ("Bluetooth: hci_bcm4377: Add new driver for BCM4377 PCIe boards")
    Fixes: 6f461f6c7c96 ("e1000e: enable/disable ASPM L0s and L1 and ERT according to
hardware errata")
    Fixes: 1eae4eb2a1c7 ("e1000e: Disable L1 ASPM power savings for 82573 mobile variants")
    Fixes: 8060e169e02f ("ath9k: Enable extended synch for AR9485 to fix L0s recovery
issue")
    Fixes: 69ce674bfa69 ("ath9k: do btcoex ASPM disabling at initialization time")
    Fixes: f37f05503575 ("mt76: mt76x2e: disable pcie_aspm by default")
    Link: https://lore.kernel.org/r/20230717120503.15276-2-ilpo.jarvinen@linux.intel.com
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: "Rafael J. Wysocki" <rafael@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 185bb783b13cc0e0eb93dc28ab5626f9901eb36c
Author: Claudiu Beznea <claudiu.beznea@microchip.com>
Date:   Wed Jun 21 13:04:09 2023 +0300

    pinctrl: mcp23s08: check return value of devm_kasprintf()

    [ Upstream commit f941714a7c7698eadb59bc27d34d6d6f38982705 ]

    devm_kasprintf() returns a pointer to dynamically allocated memory.
    Pointer could be NULL in case allocation fails. Check pointer validity.
    Identified with coccinelle (kmerr.cocci script).

    Fixes: 0f04a81784fe ("pinctrl: mcp23s08: Split to three parts: core, I²C, SPI")
    Signed-off-by: Claudiu Beznea <claudiu.beznea@microchip.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Link: https://lore.kernel.org/r/20230621100409.1608395-1-claudiu.beznea@microchip.com
    Signed-off-by: Linus Walleij <linus.walleij@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b709f83c9f79b6e2a94a8b2b865e252bcbd18a78
```

```
Author: Wu Zongyong <wuzongyong@linux.alibaba.com>
Date:   Mon Apr 10 20:34:11 2023 +0800

    PCI: Mark NVIDIA T4 GPUs to avoid bus reset

    [ Upstream commit d5af729dc2071273f14cbb94abbc60608142fd83 ]

    NVIDIA T4 GPUs do not work with SBR. This problem is found when the T4 card
    is direct attached to a Root Port only. Avoid bus reset by marking T4 GPUs
    PCI_DEV_FLAGS_NO_BUS_RESET.

    Fixes: 4c207e7121fa ("PCI: Mark some NVIDIA GPUs to avoid bus reset")
    Link:
https://lore.kernel.org/r/2dcebea53a6eb9bd212ec6d8974af2e5e0333ef6.1681129861.git.wuzongyong
@linux.alibaba.com
    Signed-off-by: Wu Zongyong <wuzongyong@linux.alibaba.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f9e8eb8ea420078d40d1e0b24a2e0589267c0d84
Author: Daire McNamara <daire.mcnamara@microchip.com>
Date:   Fri Jul 28 14:13:55 2023 +0100

    PCI: microchip: Correct the DED and SEC interrupt bit offsets

    [ Upstream commit 6d473a5a26136edf55c435a1c433e52910e03926 ]

    The SEC and DED interrupt bits are laid out the wrong way round so the SEC
    interrupt handler attempts to mask, unmask, and clear the DED interrupt
    and vice versa. Correct the bit offsets so that each interrupt handler
    operates properly.

    Link: https://lore.kernel.org/r/20230728131401.1615724-2-daire.mcnamara@microchip.com
    Fixes: 6f15a9c9f941 ("PCI: microchip: Add Microchip PolarFire PCIe controller driver")
    Signed-off-by: Daire McNamara <daire.mcnamara@microchip.com>
    Signed-off-by: Lorenzo Pieralisi <lpieralisi@kernel.org>
    Reviewed-by: Conor Dooley <conor.dooley@microchip.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aaf3420ed7cda286f22811d2b69b609fb7e803a0
Author: Luca Weiss <luca.weiss@fairphone.com>
Date:   Fri Aug 4 16:09:30 2023 +0200

    clk: qcom: gcc-sm6350: Fix gcc_sdcc2_apps_clk_src

    [ Upstream commit df04d166d1f346dbf740bbea64a3bed3e7f14c8d ]

    GPLL7 is not on by default, which causes a "gcc_sdcc2_apps_clk_src: rcg
    didn't update its configuration" error when booting. Set .flags =
    CLK_OPS_PARENT_ENABLE to fix the error.

    Fixes: 131abae905df ("clk: qcom: Add SM6350 GCC driver")
    Signed-off-by: Luca Weiss <luca.weiss@fairphone.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230804-sm6350-sdcc2-v1-1-3d946927d37d@fairphone.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 96c0cdea5b45c98935f9403e24889b4fcc16f735
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Fri Jul 28 09:57:38 2023 +0200

    clk: qcom: reset: Use the correct type of sleep/delay based on length

    [ Upstream commit 181b66ee7cdd824797fc99b53bec29cf5630a04f ]
```

```
    Use the fsleep() helper that (based on the length of the delay, see: [1])
    chooses the correct sleep/delay functions.

    [1] https://www.kernel.org/doc/Documentation/timers/timers-howto.txt

    Fixes: 2cb8a39b6781 ("clk: qcom: reset: Allow specifying custom reset delay")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230726-topic-qcom_reset-v3-1-5958facd5db2@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e487f60617178bfd7f6ba1ab19289c79c2f1be81
Author: Dmitry Torokhov <dmitry.torokhov@gmail.com>
Date:   Fri Jul 14 15:45:32 2023 -0700

    kvm/vfio: ensure kvg instance stays around in kvm_vfio_group_add()

    [ Upstream commit 9e0f4f2918c2ff145d3dedee862d9919a6ed5812 ]

    kvm_vfio_group_add() creates kvg instance, links it to kv->group_list,
    and calls kvm_vfio_file_set_kvm() with kvg->file as an argument after
    dropping kv->lock. If we race group addition and deletion calls, kvg
    instance may get freed by the time we get around to calling
    kvm_vfio_file_set_kvm().

    Previous iterations of the code did not reference kvg->file outside of
    the critical section, but used a temporary variable. Still, they had
    similar problem of the file reference being owned by kvg structure and
    potential for kvm_vfio_group_del() dropping it before
    kvm_vfio_group_add() had a chance to complete.

    Fix this by moving call to kvm_vfio_file_set_kvm() under the protection
    of kv->lock. We already call it while holding the same lock when vfio
    group is being deleted, so it should be safe here as well.

    Fixes: 2fc1bec15883 ("kvm: set/clear kvm to/from vfio_group when group add/delete")
    Reviewed-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Link: https://lore.kernel.org/r/20230714224538.404793-1-dmitry.torokhov@gmail.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de01c4c3455de19371b63abe1e79bfcc394c785f
Author: Yi Liu <yi.l.liu@intel.com>
Date:   Tue Jul 18 06:55:29 2023 -0700

    kvm/vfio: Prepare for accepting vfio device fd

    [ Upstream commit 2f99073a722beef5f74f3b0f32bda227ba3df1e0 ]

    This renames kvm_vfio_group related helpers to prepare for accepting
    vfio device fd. No functional change is intended.

    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Reviewed-by: Eric Auger <eric.auger@redhat.com>
    Reviewed-by: Jason Gunthorpe <jgg@nvidia.com>
    Tested-by: Terrence Xu <terrence.xu@intel.com>
    Tested-by: Nicolin Chen <nicolinc@nvidia.com>
    Tested-by: Matthew Rosato <mjrosato@linux.ibm.com>
    Tested-by: Yanting Jiang <yanting.jiang@intel.com>
    Tested-by: Shameer Kolothum <shameerali.kolothum.thodi@huawei.com>
    Tested-by: Zhenzhong Duan <zhenzhong.duan@intel.com>
    Signed-off-by: Yi Liu <yi.l.liu@intel.com>
```

```
    Link: https://lore.kernel.org/r/20230718135551.6592-5-yi.l.liu@intel.com
    Signed-off-by: Alex Williamson <alex.williamson@redhat.com>
    Stable-dep-of: 9e0f4f2918c2 ("kvm/vfio: ensure kvg instance stays around in
kvm_vfio_group_add()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f1b1f986d2927b5c69d7c5759518d25cabb20e56
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Tue Aug 1 12:56:32 2023 +0200

    clk: qcom: fix some Kconfig corner cases

    [ Upstream commit b6bcd1c0c27e1f210228346e6d23a2ec0c263e8c ]

    The SM_GCC_8550 symbol and others can only be built for ARM64 or when
    compile testing, but it gets selected by other drivers that can also be
    built for 32-bit ARCH_QCOM when not compile testing, which results in
    a Kconfig warning:

    WARNING: unmet direct dependencies detected for SM_GCC_8550
      Depends on [n]: COMMON_CLK [=y] && COMMON_CLK_QCOM [=m] && (ARM64 || COMPILE_TEST
[=n])
      Selected by [m]:
      - SM_GPUCC_8550 [=m] && COMMON_CLK [=y] && COMMON_CLK_QCOM [=m]
      - SM_VIDEOCC_8550 [=m] && COMMON_CLK [=y] && COMMON_CLK_QCOM [=m]

    Add further 'depends on' statements to tighten this in a way that
    avoids the missing dependencies.

    Fixes: fd0b5b106fcab ("clk: qcom: Introduce SM8350 VIDEOCC")
    Fixes: 441fe711be384 ("clk: qcom: videocc-sm8450: Add video clock controller driver for
SM8450")
    Fixes: f53153a37969c ("clk: qcom: videocc-sm8550: Add video clock controller driver for
SM8550")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230801105718.3658612-1-arnd@kernel.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd7751dd7d930eaa3fd67169c7c7c79a68124e50
Author: Patrick Whewell <patrick.whewell@sightlineapplications.com>
Date:   Wed Aug 2 14:04:00 2023 -0700

    clk: qcom: gcc-sm8250: Fix gcc_sdcc2_apps_clk_src

    [ Upstream commit 783cb693828ce487cf0bc6ad16cbcf2caae6f8d9 ]

    GPLL9 is not on by default, which causes a "gcc_sdcc2_apps_clk_src: rcg
    didn't update its configuration" error when booting. Set .flags =
    CLK_OPS_PARENT_ENABLE to fix the error.

    Fixes: 3e5770921a88 ("clk: qcom: gcc: Add global clock controller driver for SM8250")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Signed-off-by: Patrick Whewell <patrick.whewell@sightlineapplications.com>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Link: https://lore.kernel.org/r/20230802210359.408-1-
patrick.whewell@sightlineapplications.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3949327e37d97b3163ac4319fa0caca7251a5924
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Tue Aug 1 22:31:56 2023 +0800
```

```
    ext4: avoid potential data overflow in next_linear_group

    [ Upstream commit 60c672b7f2d1e5dd1774f2399b355c9314e709f8 ]

    ngroups is ext4_group_t (unsigned int) while next_linear_group treat it
    in int. If ngroups is bigger than max number described by int, it will
    be treat as a negative number. Then "return group + 1 >= ngroups ? 0 :
    group + 1;" may keep returning 0.
    Switch int to ext4_group_t in next_linear_group to fix the overflow.

    Fixes: 196e402adf2e ("ext4: improve cr 0 / cr 1 group scanning")
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Reviewed-by: Ritesh Harjani (IBM) <ritesh.list@gmail.com>
    Link: https://lore.kernel.org/r/20230801143204.2284343-3-shikemeng@huaweicloud.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e69d665987db0e37896adf78a7e718f9a0a75d3f
Author: Kemeng Shi <shikemeng@huaweicloud.com>
Date:   Tue Aug 1 22:31:55 2023 +0800

    ext4: correct grp validation in ext4_mb_good_group

    [ Upstream commit a9ce5993a0f5c0887c8a1b4ffa3b8046fbcfdc93 ]

    Group corruption check will access memory of grp and will trigger kernel
    crash if grp is NULL. So do NULL check before corruption check.

    Fixes: 5354b2af3406 ("ext4: allow ext4_get_group_info() to fail")
    Signed-off-by: Kemeng Shi <shikemeng@huaweicloud.com>
    Reviewed-by: Ritesh Harjani (IBM) <ritesh.list@gmail.com>
    Link: https://lore.kernel.org/r/20230801143204.2284343-2-shikemeng@huaweicloud.com
    Signed-off-by: Theodore Ts'o <tytso@mit.edu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5390840eb917e9abf060de594ffcfb5054add8db
Author: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
Date:   Tue Jul 25 16:04:27 2023 +0800

    EDAC/igen6: Fix the issue of no error events

    [ Upstream commit ce53ad81ed36c24aff075f94474adecfabfcf239 ]

    Current igen6_edac checks for pending errors before the registration
    of the error handler. However, there is a possibility that the error
    occurs during the registration process, leading to unhandled pending
    errors and no future error events. This issue can be reproduced by
    repeatedly injecting errors during the loading of the igen6_edac.

    Fix this issue by moving the pending error handler after the registration
    of the error handler, ensuring that no pending errors are left unhandled.

    Fixes: 10590a9d4f23 ("EDAC/igen6: Add EDAC driver for Intel client SoCs using IBECC")
    Reported-by: Ee Wey Lim <ee.wey.lim@intel.com>
    Tested-by: Ee Wey Lim <ee.wey.lim@intel.com>
    Signed-off-by: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
    Signed-off-by: Tony Luck <tony.luck@intel.com>
    Link: https://lore.kernel.org/r/20230725080427.23883-1-qiuxu.zhuo@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8d1bf997d8e602a1993a611d1e89ffd437abcf3e
Author: David Wronek <davidwronek@gmail.com>
Date:   Sun Jul 23 21:05:02 2023 +0200
```

```
    clk: qcom: gcc-sc7180: Fix up gcc_sdcc2_apps_clk_src

    [ Upstream commit fd0b5ba87ad5709f0fd3d2bc4b7870494a75f96a ]

    Set .flags = CLK_OPS_PARENT_ENABLE to fix "gcc_sdcc2_apps_clk_src: rcg
    didn't update its configuration" error.

    Fixes: 17269568f726 ("clk: qcom: Add Global Clock controller (GCC) driver for SC7180")
    Signed-off-by: David Wronek <davidwronek@gmail.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230723190725.1619193-2-davidwronek@gmail.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f22b7c9359699a1fad4cf7a39c4aab99707a3bbe
Author: Zhang Jianhua <chris.zjh@huawei.com>
Date:   Sat Jul 22 15:31:07 2023 +0000

    clk: sunxi-ng: Modify mismatched function name

    [ Upstream commit 075d9ca5b4e17f84fd1c744a405e69ec743be7f0 ]

    No functional modification involved.

    drivers/clk/sunxi-ng/ccu_mmc_timing.c:54: warning: expecting prototype for
sunxi_ccu_set_mmc_timing_mode(). Prototype was for sunxi_ccu_get_mmc_timing_mode() instead

    Fixes: f6f64ed868d3 ("clk: sunxi-ng: Add interface to query or configure MMC timing
modes.")
    Signed-off-by: Zhang Jianhua <chris.zjh@huawei.com>
    Reviewed-by: Randy Dunlap <rdunlap@infradead.org>
    Link: https://lore.kernel.org/r/20230722153107.2078179-1-chris.zjh@huawei.com
    Signed-off-by: Jernej Skrabec <jernej.skrabec@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9843bdd8bbad98d822046e4a551e4cb44dee348c
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jul 25 10:51:56 2023 +0200

    clk: qcom: dispcc-sc8280xp: Use ret registers on GDSCs

    [ Upstream commit 20e1d75bc043c5ec1fd8f5169fde17db89eb11c3 ]

    The DISP_CC GDSCs have not been instructed to use the ret registers.
    Fix that.

    Fixes: 4a66e76fdb6d ("clk: qcom: Add SC8280XP display clock controller")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230725-topic-8280_dispcc_gdsc-v1-1-
236590060531@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 19cf3ba16dcc2ef059dcf010072d4f96d76486e0
Author: Ira Weiny <ira.weiny@intel.com>
Date:   Wed Jul 26 11:29:42 2023 -0700

    PCI/DOE: Fix destroy_work_on_stack() race

    [ Upstream commit e3a3a097eaebaf234a482b4d2f9f18fe989208c1 ]

    The following debug object splat was observed in testing:

      ODEBUG: free active (active state 0) object: 0000000097d23782 object type: work_struct
hint: doe_statemachine_work+0x0/0x510
```

```
       WARNING: CPU: 1 PID: 71 at lib/debugobjects.c:514 debug_print_object+0x7d/0xb0
       ...
       Workqueue: pci 0000:36:00.0 DOE [1 doe_statemachine_work
       RIP: 0010:debug_print_object+0x7d/0xb0
       ...
       Call Trace:
        ? debug_print_object+0x7d/0xb0
        ? __pfx_doe_statemachine_work+0x10/0x10
        debug_object_free.part.0+0x11b/0x150
        doe_statemachine_work+0x45e/0x510
        process_one_work+0x1d4/0x3c0

    This occurs because destroy_work_on_stack() was called after signaling
    the completion in the calling thread.  This creates a race between
    destroy_work_on_stack() and the task->work struct going out of scope in
    pci_doe().

    Signal the work complete after destroying the work struct.  This is safe
    because signal_task_complete() is the final thing the work item does and
    the workqueue code is careful not to access the work struct after.

    Fixes: abf04be0e707 ("PCI/DOE: Fix memory leak with CONFIG_DEBUG_OBJECTS=y")
    Link: https://lore.kernel.org/r/20230726-doe-fix-v1-1-af07e614d4dd@intel.com
    Signed-off-by: Ira Weiny <ira.weiny@intel.com>
    Signed-off-by: Bjorn Helgaas <bhelgaas@google.com>
    Reviewed-by: Lukas Wunner <lukas@wunner.de>
    Acked-by: Dan Williams <dan.j.williams@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4922517fd5631f508d78216225119a530fd26a2e
Author: Jason Gunthorpe <jgg@ziepe.ca>
Date:   Mon Jul 17 15:12:07 2023 -0300

    iommufd: Fix locking around hwpt allocation

    [ Upstream commit 31422dff187b243c58f3a97d16bbe9e9ada639fe ]

    Due to the auto_domains mechanism the ioas->mutex must be held until
    the hwpt is completely setup by iommufd_object_abort_and_destroy() or
    iommufd_object_finalize().

    This prevents a concurrent iommufd_device_auto_get_domain() from seeing
    an incompletely initialized object through the ioas->hwpt_list.

    To make this more consistent move the unlock until after finalize.

    Fixes: e8d57210035b ("iommufd: Add kAPI toward external drivers for physical devices")
    Link: https://lore.kernel.org/r/11-v8-6659224517ea+532-iommufd_alloc_jgg@nvidia.com
    Reviewed-by: Kevin Tian <kevin.tian@intel.com>
    Tested-by: Nicolin Chen <nicolinc@nvidia.com>
    Signed-off-by: Jason Gunthorpe <jgg@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e983d11482af44b3c7b7cf67bc2ee2bc0bcc62a2
Author: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
Date:   Mon Jul 10 09:32:32 2023 +0800

    EDAC/i10nm: Skip the absent memory controllers

    [ Upstream commit c545f5e412250555bd4e717d062b117f20bab418 ]

    Some Sapphire Rapids workstations' absent memory controllers
    still appear as PCIe devices that fool the i10nm_edac driver
    and result in "shift exponent -66 is negative" call traces
    from skx_get_dimm_info().
```

```
    Skip the absent memory controllers to avoid the call traces.

    Reported-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Closes: https://lore.kernel.org/linux-
edac/CAAd53p41Ku1m1rapeqb1xtD+kKuk+BaUW=dumuoF0ZO3GhFjFA@mail.gmail.com/T/#m5de16dce60a8c836
ec235868c7c16e3fefad0cc2
    Tested-by: Kai-Heng Feng <kai.heng.feng@canonical.com>
    Reported-by: Koba Ko <koba.ko@canonical.com>
    Closes: https://lore.kernel.org/linux-
edac/SA1PR11MB71305B71CCCC3D9305835202892AA@SA1PR11MB7130.namprd11.prod.outlook.com/T/#t
    Tested-by: Koba Ko <koba.ko@canonical.com>
    Fixes: d4dc89d069aa ("EDAC, i10nm: Add a driver for Intel 10nm server processors")
    Signed-off-by: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
    Signed-off-by: Tony Luck <tony.luck@intel.com>
    Link: https://lore.kernel.org/r/20230710013232.59712-1-qiuxu.zhuo@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ad428064306cbf64426d9df5b8f016f78ac6ec1a
Author: Minjie Du <duminjie@vivo.com>
Date:   Wed Jul 12 18:22:46 2023 +0800

    drivers: clk: keystone: Fix parameter judgment in _of_pll_clk_init()

    [ Upstream commit a995c50db887ef97f3160775aef7d772635a6f6e ]

    The function clk_register_pll() may return NULL or an ERR_PTR. Don't
    treat an ERR_PTR as valid.

    Signed-off-by: Minjie Du <duminjie@vivo.com>
    Link: https://lore.kernel.org/r/20230712102246.10348-1-duminjie@vivo.com
    Fixes: b9e0d40c0d83 ("clk: keystone: add Keystone PLL clock driver")
    [sboyd@kernel.org: Reword commit text]
    Signed-off-by: Stephen Boyd <sboyd@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8aed38fa6150055934dfa3aab8403fad5691fcfd
Author: Johan Hovold <johan+linaro@kernel.org>
Date:   Tue Jul 18 15:28:58 2023 +0200

    clk: qcom: gcc-sc8280xp: fix runtime PM imbalance on probe errors

    [ Upstream commit 10192ab375c39c58d39cba028d9685cefe1ca3c2 ]

    Make sure to decrement the runtime PM usage count before returning in
    case RCG dynamic frequency switch initialisation fails.

    Fixes: 2a541abd9837 ("clk: qcom: gcc-sc8280xp: Add runtime PM")
    Cc: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Johan Hovold <johan+linaro@kernel.org>
    Link: https://lore.kernel.org/r/20230718132902.21430-5-johan+linaro@kernel.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15d9911f33d4fd87d2ffcf39851363587ef9d8b8
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Tue Jun 27 19:40:36 2023 +0530

    PCI: qcom-ep: Switch MHI bus master clock off during L1SS

    [ Upstream commit b9cbc06049cb6b7a322d708c2098195fb9fdcc4c ]

    Currently, as part of the qcom_pcie_perst_deassert() function, instead
    of writing the updated value to clear PARF_MSTR_AXI_CLK_EN, the variable
    "val" is re-read.
```

```
    This must be fixed to ensure that the master clock supplied to the MHI
    bus is correctly gated during L1.1/L1.2 to save power.

    Thus, replace the line that re-reads "val" with a line that writes the
    updated value to the register to clear PARF_MSTR_AXI_CLK_EN.

    [kwilczynski: commit log]
    Fixes: c457ac029e44 ("PCI: qcom-ep: Gate Master AXI clock to MHI bus during L1SS")
    Link: https://lore.kernel.org/linux-pci/20230627141036.11600-1-
manivannan.sadhasivam@linaro.org
    Reported-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7992c946c846e8b1961855e8e6dc7184e5445bd1
Author: Sven Peter <sven@svenpeter.dev>
Date:   Sat Mar 11 14:34:53 2023 +0100

    PCI: apple: Initialize pcie->nvecs before use

    [ Upstream commit d8650c0c2aa2e413594e4cb0faafa9958c1d7782 ]

    The apple_pcie_setup_port() function computes ilog2(pcie->nvecs) to set
    up the number of MSIs available for each port. However, it's called
    before apple_msi_init(), which initializes pcie->nvecs.

    Luckily, pcie->nvecs is part of kzalloc()-ed structure and, as such,
    initialized as zero. ilog2(0) happens to be 0xffffffff which then simply
    configures more MSIs in hardware than we have. This doesn't break
    anything because we never hand out those vectors.

    Thus, swap the order of the two calls so that the correctly initialized
    value is then used.

    [kwilczynski: commit log]
    Link: https://lore.kernel.org/linux-pci/20230311133453.63246-1-sven@svenpeter.dev
    Fixes: 476c41ed4597 ("PCI: apple: Implement MSI support")
    Signed-off-by: Sven Peter <sven@svenpeter.dev>
    Signed-off-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Reviewed-by: Marc Zyngier <maz@kernel.org>
    Reviewed-by: Alyssa Rosenzweig <alyssa@rosenzweig.io>
    Reviewed-by: Eric Curtin <ecurtin@redhat.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d0c8d2e9e51a9a9502a1b236dfffbef8ff5925d8
Author: Alibek Omarov <a1ba.omarov@gmail.com>
Date:   Wed Jun 14 16:47:50 2023 +0300

    clk: rockchip: rk3568: Fix PLL rate setting for 78.75MHz

    [ Upstream commit dafebd0f9a4f56b10d7fbda0bff1f540d16a2ea4 ]

    PLL rate on RK356x is calculated through the simple formula:
    ((24000000 / _refdiv) * _fbdiv) / (_postdiv1 * _postdiv2)

    The PLL rate setting for 78.75MHz seems to be copied from 96MHz
    so this patch fixes it and configures it properly.

    Signed-off-by: Alibek Omarov <a1ba.omarov@gmail.com>
    Fixes: 842f4cb72639 ("clk: rockchip: Add more PLL rates for rk3568")
    Reviewed-by: Sascha Hauer <s.hauer@pengutronix.de>
    Link: https://lore.kernel.org/r/20230614134750.1056293-1-a1ba.omarov@gmail.com
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
```

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

commit 539201e051335023dd752a8bd78b593e7a8ceb10
Author: Danila Tikhonov <danila@jiaxyga.com>
Date:   Fri Jun 30 22:19:44 2023 +0300

    clk: qcom: gcc-sm7150: Add CLK_OPS_PARENT_ENABLE to sdcc2 rcg

    [ Upstream commit ff19022b9112d6bbd7c117c83e944cb21b438e91 ]

    Set .flags = CLK_OPS_PARENT_ENABLE to fix "gcc_sdcc2_apps_clk_src: rcg
    didn't update its configuration" error.

    Fixes: a808d58ddf29 ("clk: qcom: Add Global Clock Controller (GCC) driver for SM7150")
    Signed-off-by: Danila Tikhonov <danila@jiaxyga.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230630191944.20282-1-danila@jiaxyga.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a92a9604e8a431b92ff70a952f346099033bf2f4
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 19:48:08 2023 +0200

    clk: qcom: gcc-sc8280xp: Add missing GDSCs

    [ Upstream commit 4712eb7ff85bd3dd09c6668b8de4080e02b3eea9 ]

    There are 10 more GDSCs that we've not been caring about, and by extension
    (and perhaps even more importantly), not putting to sleep. Add them.

    Fixes: d65d005f9a6c ("clk: qcom: add sc8280xp GCC driver")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Acked-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20230620-topic-sc8280_gccgdsc-v2-3-
562c1428c10d@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 66120ba55999a33ba1b259f521b7f2e482b72f10
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 19:48:07 2023 +0200

    dt-bindings: clock: qcom,gcc-sc8280xp: Add missing GDSCs

    [ Upstream commit 9eba4db02a88e7a810aabd70f7a6960f184f391f ]

    There are 10 more GDSCs that we've not been caring about, and by extension
    (and perhaps even more importantly), not putting to sleep. Add them.

    Fixes: a66a82f2a55e ("dt-bindings: clock: Add Qualcomm SC8280XP GCC bindings")
    Acked-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Acked-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20230620-topic-sc8280_gccgdsc-v2-2-
562c1428c10d@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Stable-dep-of: 4712eb7ff85b ("clk: qcom: gcc-sc8280xp: Add missing GDSCs")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a5828a10686e3f339e4d639fb9ac5282653b83ff
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 19:48:06 2023 +0200

    clk: qcom: gcc-sc8280xp: Add missing GDSC flags
```

```
    [ Upstream commit 2fd02de27054576a4a8c89302e2f77122c55e957 ]

    All of the 8280's GCC GDSCs can and should use the retain registers so
    as not to lose their state when entering lower power modes.

    Fixes: d65d005f9a6c ("clk: qcom: add sc8280xp GCC driver")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Acked-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20230620-topic-sc8280_gccgdsc-v2-1-
562c1428c10d@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 254971315e544e0453fc9890bd2be2e2d49c4461
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Jun 14 13:35:33 2023 +0200

    clk: qcom: gpucc-sm6350: Fix clock source names

    [ Upstream commit 743913b343a3ec2510fe3c0dfaff03d049659922 ]

    fw_name for GCC inputs didn't match the bindings. Fix it.

    Fixes: 013804a727a0 ("clk: qcom: Add GPU clock controller driver for SM6350")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230315-topic-lagoon_gpu-v2-2-afcdfb18bb13@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3f85d8fdcbdf5a0ff1b711bc6a25016fb44dc147
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Jun 14 13:35:32 2023 +0200

    clk: qcom: gpucc-sm6350: Introduce index-based clk lookup

    [ Upstream commit f6f89d194e4ddcfe197ac8a05ed4161f642a5c68 ]

    Add the nowadays-prefered and marginally faster way of looking up parent
    clocks in the device tree. It also allows for clock-names-independent
    operation, so long as the order (which is enforced by schema) is kept.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230315-topic-lagoon_gpu-v2-1-afcdfb18bb13@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Stable-dep-of: 743913b343a3 ("clk: qcom: gpucc-sm6350: Fix clock source names")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b870caeb18041f856893066ded81c560db3d56cc
Author: Corey Minyard <minyard@acm.org>
Date:   Mon Jun 19 11:43:33 2023 -0500

    ipmi:ssif: Fix a memory leak when scanning for an adapter

    [ Upstream commit b8d72e32e1453d37ee5c8a219f24e7eeadc471ef ]

    The adapter scan ssif_info_find() sets info->adapter_name if the adapter
    info came from SMBIOS, as it's not set in that case.  However, this
    function can be called more than once, and it will leak the adapter name
    if it had already been set.  So check for NULL before setting it.

    Fixes: c4436c9149c5 ("ipmi_ssif: avoid registering duplicate ssif interface")
    Signed-off-by: Corey Minyard <minyard@acm.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit dddba671848570f5607c20c308d7fc9913826f72
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Mon Jun 19 17:28:02 2023 +0800

    ipmi:ssif: Add check for kstrdup

    [ Upstream commit c5586d0f711e9744d0cade39b0c4a2d116a333ca ]

    Add check for the return value of kstrdup() and return the error
    if it fails in order to avoid NULL pointer dereference.

    Fixes: c4436c9149c5 ("ipmi_ssif: avoid registering duplicate ssif interface")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Message-Id: <20230619092802.35384-1-jiasheng@iscas.ac.cn>
    Signed-off-by: Corey Minyard <minyard@acm.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb0e6ba256acbc8dc09f835eec6c7f21827ad76f
Author: Takashi Iwai <tiwai@suse.de>
Date:   Sat Aug 26 09:21:51 2023 +0200

    ALSA: ump: Fix -Wformat-truncation warnings

    [ Upstream commit 4aa69d64e43edb51a4ecff7d301e9f881eb2d3f5 ]

    Filling the rawmidi name and substream name can be truncated, and this
    leads to spurious compiler warnings due to -Wformat-truncation.
    Although the truncation is the expected behavior, it'd be better to
    truncate the string within "(...)"

    This patch puts the precision specifies to each %s for fitting the
    words within the size-limited strings.

    Fixes: 5f11dd938fe7 ("ALSA: usb-audio: Attach legacy rawmidi after probing all UMP EPs")
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202308251844.1FuQYsql-lkp@intel.com/
    Link: https://lore.kernel.org/r/20230826072151.23408-1-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fe4fa3352661c8e0234f2578be455aaadeaa01d
Author: Takashi Iwai <tiwai@suse.de>
Date:   Thu Aug 24 09:51:07 2023 +0200

    ALSA: ump: Don't create unused substreams for static blocks

    [ Upstream commit b2bcbd031d34d1ba1f491b9152474cf9f6d4d51b ]

    When the UMP Endpoint is declared as "static", that is, no dynamic
    reassignment of UMP Groups, it makes little sense to expose always all
    16 groups with 16 substreams.  Many of those substreams are disabled
    groups, hence they are useless, but applications don't know it and try
    to open / access all those substreams unnecessarily.

    This patch limits the number of UMP legacy rawmidi substreams only to
    the active groups.  The behavior is changed only for the static
    endpoint (i.e. devices without UMP v1.1 feature implemented or with
    the static block flag is set).

    Fixes: 0b5288f5fe63 ("ALSA: ump: Add legacy raw MIDI support")
    Link: https://lore.kernel.org/r/20230824075108.29958-4-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ab66172e48ef0fd7595574d4bd25c8db063f950a
```

```
Author: Takashi Iwai <tiwai@suse.de>
Date:   Thu Aug 24 09:51:06 2023 +0200

    ALSA: ump: Fill group names for legacy rawmidi substreams

    [ Upstream commit 1761f4cc114af531020ea190df6a24dd288a8221 ]

    To make it clearer which legacy substream corresponds to which UMP
    group, fill the subname field of each substream object with the group
    number and the endpoint name, e.g. "Group 1 (My Device)".

    Ideally speaking, we should have some better link information to the
    derived UMP, but it's another feature extension.

    Fixes: 0b5288f5fe63 ("ALSA: ump: Add legacy raw MIDI support")
    Link: https://lore.kernel.org/r/20230824075108.29958-3-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5cb4f6e220a8b02cc957af2f5fe2ba295719a87b
Author: Takashi Iwai <tiwai@suse.de>
Date:   Thu Aug 24 09:51:05 2023 +0200

    ALSA: usb-audio: Attach legacy rawmidi after probing all UMP EPs

    [ Upstream commit 5f11dd938fe7657899ca79b2ffc4d708e43f4737 ]

    The legacy rawmidi devices are the shadows of the main UMP devices,
    hence it's better to initialize them after all UMP Endpoints are
    parsed.  Then, at the moment the legacy rawmidi is created, we already
    know the static flag or the proper EP name string, and we can fill
    those information at UMP core side instead of fiddling the attributes
    at a later point.

    Fixes: ec362b63c4b5 ("ALSA: usb-audio: Enable the legacy raw MIDI support")
    Link: https://lore.kernel.org/r/20230824075108.29958-2-tiwai@suse.de
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9504deaf0c6d2928815a7377db5fe6fecc0b8272
Author: Su Hui <suhui@nfschina.com>
Date:   Wed Aug 23 10:52:13 2023 +0800

    ALSA: ac97: Fix possible error value of *rac97

    [ Upstream commit 67de40c9df94037769967ba28c7d951afb45b7fb ]

    Before committing 79597c8bf64c, *rac97 always be NULL if there is
    an error. When error happens, make sure *rac97 is NULL is safer.

    For examble, in snd_vortex_mixer():
            err = snd_ac97_mixer(pbus, &ac97, &vortex->codec);
            vortex->isquad = ((vortex->codec == NULL) ?
                    0 : (vortex->codec->ext_id&0x80));
    If error happened but vortex->codec isn't NULL, this may cause some
    problems.

    Move the judgement order to be clearer and better.

    Fixes: 79597c8bf64c ("ALSA: ac97: Fix possible NULL dereference in snd_ac97_mixer")
    Suggested-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Acked-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Signed-off-by: Su Hui <suhui@nfschina.com>
    Link: https://lore.kernel.org/r/20230823025212.1000961-1-suhui@nfschina.com
    Signed-off-by: Takashi Iwai <tiwai@suse.de>
```

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

commit b2673e3af576c421aa1ee08bb6a35d2c062d430f
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Jul 28 10:50:29 2023 +0200

    of: unittest: Fix overlay type in apply/revert check

    [ Upstream commit 6becf8f845ae1f0b1cfed395bbeccbd23654162d ]

    The removal check in of_unittest_apply_revert_overlay_check()
    always uses the platform device overlay type, while it should use the
    actual overlay type, as passed as a parameter to the function.

    This has no impact on any current test, as all tests calling
    of_unittest_apply_revert_overlay_check() use the platform device overlay
    type.

    Fixes: d5e75500ca401d31 ("of: unitest: Add I2C overlay unit tests.")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/r/ba0234c41ba808f10112094f88792beeb6dbaedf.1690533838.git.geert+rene
sas@glider.be
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c403c81b577a67fe9ec6a2e89d143256487be50f
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Fri Jul 28 10:50:28 2023 +0200

    of: overlay: Call of_changeset_init() early

    [ Upstream commit a9515ff4fb142b690a0d2b58782b15903b990dba ]

    When of_overlay_fdt_apply() fails, the changeset may be partially
    applied, and the caller is still expected to call of_overlay_remove() to
    clean up this partial state.

    However, of_overlay_apply() calls of_resolve_phandles() before
    init_overlay_changeset().  Hence if the overlay fails to apply due to an
    unresolved symbol, the overlay_changeset.cset.entries list is still
    uninitialized, and cleanup will crash with a NULL-pointer dereference in
    overlay_removal_is_ok().

    Fix this by moving the call to of_changeset_init() from
    init_overlay_changeset() to of_overlay_fdt_apply(), where all other
    early initialization is done.

    Fixes: f948d6d8b792bb90 ("of: overlay: avoid race condition between applying multiple
overlays")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Link:
https://lore.kernel.org/r/4f1d6d74b61cba2599026adb6d1948ae559ce91f.1690533838.git.geert+rene
sas@glider.be
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6c23744a3e960f207843e7ca17c1241c59ff470e
Author: Vijendar Mukunda <Vijendar.Mukunda@amd.com>
Date:   Wed Aug 23 13:03:39 2023 +0530

    ASoC: SOF: amd: clear dsp to host interrupt status

    [ Upstream commit 38592ae6dc9f84b7a994c43de2136b8115ca30f6 ]
```

```
    DSP_SW_INTR_STAT_OFFSET is a common interrupt register which will be
    accessed by both ACP firmware and driver. This register contains register
    bits corresponds to host to dsp interrupts and vice versa.

    when dsp to host interrupt is reported, only clear dsp to host
    interrupt bit in DSP_SW_INTR_STAT_OFFSET.

    Fixes: 2e7c6652f9b8 ("ASoC: SOF: amd: Fix for handling spurious interrupts from DSP")

    Signed-off-by: Vijendar Mukunda <Vijendar.Mukunda@amd.com>
    Link: https://lore.kernel.org/r/20230823073340.2829821-7-Vijendar.Mukunda@amd.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4e142a3433884543fe62ed174c0fac31aa5dd37c
Author: David Jeffery <djeffery@redhat.com>
Date:   Wed Aug 16 14:13:55 2023 -0400

    md: raid0: account for split bio in iostat accounting

    [ Upstream commit cc22b5407e9ca76adb7efeed843146510b1b72a5 ]

    When a bio is split by md raid0, the newly created bio will not be tracked
    by md for I/O accounting. Only the portion of I/O still assigned to the
    original bio which was reduced by the split will be accounted for. This
    results in md iostat data sometimes showing I/O values far below the actual
    amount of data being sent through md.

    md_account_bio() needs to be called for all bio generated by the bio split.

    A simple example of the issue was generated using a raid0 device on partitions
    to the same device. Since all raid0 I/O then goes to one device, it makes it
    easy to see a gap between the md device and its sd storage. Reading an lvm
    device on top of the md device, the iostat output (some 0 columns and extra
    devices removed to make the data more compact) was:

    Device             tps    kB_read/s    kB_wrtn/s    kB_dscd/s    kB_read
    md2               0.00         0.00         0.00         0.00          0
    sde               0.00         0.00         0.00         0.00          0
    md2            1364.00    411496.00         0.00         0.00     411496
    sde            1734.00    646144.00         0.00         0.00     646144
    md2            1699.00    510680.00         0.00         0.00     510680
    sde            2155.00    802784.00         0.00         0.00     802784
    md2             803.00    241480.00         0.00         0.00     241480
    sde            1016.00    377888.00         0.00         0.00     377888
    md2               0.00         0.00         0.00         0.00          0
    sde               0.00         0.00         0.00         0.00          0

    I/O was generated doing large direct I/O reads (12M) with dd to a linear
    lvm volume on top of the 4 leg raid0 device.

    The md2 reads were showing as roughly 2/3 of the reads to the sde device
    containing all of md2's raid partitions. The sum of reads to sde was
    1826816 kB, which was the expected amount as it was the amount read by
    dd. With the patch, the total reads from md will match the reads from
    sde and be consistent with the amount of I/O generated.

    Fixes: 10764815ff47 ("md: add io accounting for raid0 and raid5")
    Signed-off-by: David Jeffery <djeffery@redhat.com>
    Tested-by: Laurence Oberman <loberman@redhat.com>
    Reviewed-by: Laurence Oberman <loberman@redhat.com>
    Reviewed-by: Yu Kuai <yukuai3@huawei.com>
    Signed-off-by: Song Liu <song@kernel.org>
    Link: https://lore.kernel.org/r/20230816181433.13289-1-djeffery@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit ae4abf22c41f19e7685d9674f05d26ef088448b4
Author: Jan Kara <jack@suse.cz>
Date:   Mon Aug 14 11:27:08 2023 +0200

    md/raid0: Fix performance regression for large sequential writes

    [ Upstream commit 319ff40a542736d67e5bce18635de35d0e7a0bff ]

    Commit f00d7c85be9e ("md/raid0: fix up bio splitting.") among other
    things changed how bio that needs to be split is submitted. Before this
    commit, we have split the bio, mapped and submitted each part. After
    this commit, we map only the first part of the split bio and submit the
    second part unmapped. Due to bio sorting in __submit_bio_noacct() this
    results in the following request ordering:

      9,0   18     1181     0.525037895 15995  Q  WS 1479315464 + 63392

      Split off chunk-sized (1024 sectors) request:

      9,0   18     1182     0.629019647 15995  X  WS 1479315464 / 1479316488

      Request is unaligned to the chunk so it's split in
      raid0_make_request().  This is the first part mapped and punted to
      bio_list:

      8,0   18     7053     0.629020455 15995  A  WS 739921928 + 1016 <- (9,0) 1479315464

      Now raid0_make_request() returns, second part is postponed on
      bio_list. __submit_bio_noacct() resorts the bio_list, mapped request
      is submitted to the underlying device:

      8,0   18     7054     0.629022782 15995  G  WS 739921928 + 1016

      Now we take another request from the bio_list which is the remainder
      of the original huge request. Split off another chunk-sized bit from
      it and the situation repeats:

      9,0   18     1183     0.629024499 15995  X  WS 1479316488 / 1479317512
      8,16  18     6998     0.629025110 15995  A  WS 739921928 + 1016 <- (9,0) 1479316488
      8,16  18     6999     0.629026728 15995  G  WS 739921928 + 1016
      ...
      9,0   18     1184     0.629032940 15995  X  WS 1479317512 / 1479318536 [libnetacq-
write]
      8,0   18     7059     0.629033294 15995  A  WS 739922952 + 1016 <- (9,0) 1479317512
      8,0   18     7060     0.629033902 15995  G  WS 739922952 + 1016
      ...

      This repeats until we consume the whole original huge request. Now we
      finally get to processing the second parts of the split off requests
      (in reverse order):

      8,16  18     7181     0.629161384 15995  A  WS 739952640 + 8 <- (9,0) 1479377920
      8,0   18     7239     0.629162140 15995  A  WS 739952640 + 8 <- (9,0) 1479376896
      8,16  18     7186     0.629163881 15995  A  WS 739951616 + 8 <- (9,0) 1479375872
      8,0   18     7242     0.629164421 15995  A  WS 739951616 + 8 <- (9,0) 1479374848
      ...

    I guess it is obvious that this IO pattern is extremely inefficient way
    to perform sequential IO. It also makes bio_list to grow to rather long
    lengths.

    Change raid0_make_request() to map both parts of the split bio. Since we
    know we are provided with at most chunk-sized bios, we will always need
    to split the incoming bio at most once.
```

```
    Fixes: f00d7c85be9e ("md/raid0: fix up bio splitting.")
    Signed-off-by: Jan Kara <jack@suse.cz>
    Reviewed-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230814092720.3931-2-jack@suse.cz
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50c1e6664962fca0b09d668c0b4ea114d84e2698
Author: Jan Kara <jack@suse.cz>
Date:   Mon Aug 14 11:27:07 2023 +0200

    md/raid0: Factor out helper for mapping and submitting a bio

    [ Upstream commit af50e20afb401cc203bd2a9ff62ece0ae4976103 ]

    Factor out helper function for mapping and submitting a bio out of
    raid0_make_request(). We will use it later for submitting both parts of
    a split bio.

    Signed-off-by: Jan Kara <jack@suse.cz>
    Reviewed-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230814092720.3931-1-jack@suse.cz
    Signed-off-by: Song Liu <song@kernel.org>
    Stable-dep-of: 319ff40a5427 ("md/raid0: Fix performance regression for large sequential
writes")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d63c6f9765339dcfc34b7365ced7c518012e4fe
Author: Tejun Heo <tj@kernel.org>
Date:   Wed Aug 16 09:56:23 2023 -1000

    blk-cgroup: Fix NULL deref caused by blkg_policy_data being installed before init

    [ Upstream commit ec14a87ee1999b19d8b7ed0fa95fea80644624ae ]

    blk-iocost sometimes causes the following crash:

      BUG: kernel NULL pointer dereference, address: 00000000000000e0
      ...
      RIP: 0010:_raw_spin_lock+0x17/0x30
      Code: be 01 02 00 00 e8 79 38 39 ff 31 d2 89 d0 5d c3 0f 1f 00 0f 1f 44 00 00 55 48 89
e5 65 ff 05 48 d0 34 7e b9 01 00 00 00 31 c0 <f0> 0f b1 0f 75 02 5d c3 89 c6 e8 ea 04 00 00
5d c3 0f 1f 84 00 00
      RSP: 0018:ffffc900023b3d40 EFLAGS: 00010046
      RAX: 0000000000000000 RBX: 00000000000000e0 RCX: 0000000000000001
      RDX: ffffc900023b3d20 RSI: ffffc900023b3cf0 RDI: 00000000000000e0
      RBP: ffffc900023b3d40 R08: ffffc900023b3c10 R09: 0000000000000003
      R10: 0000000000000064 R11: 000000000000000a R12: ffff888102337000
      R13: fffffffffffffff2 R14: ffff88810af408c8 R15: ffff8881070c3600
      FS:  00007faaaf364fc0(0000) GS:ffff88842fdc0000(0000) knlGS:0000000000000000
      CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
      CR2: 00000000000000e0 CR3: 00000001097b1000 CR4: 0000000000350ea0
      Call Trace:
       <TASK>
       ioc_weight_write+0x13d/0x410
       cgroup_file_write+0x7a/0x130
       kernfs_fop_write_iter+0xf5/0x170
       vfs_write+0x298/0x370
       ksys_write+0x5f/0xb0
       __x64_sys_write+0x1b/0x20
       do_syscall_64+0x3d/0x80
       entry_SYSCALL_64_after_hwframe+0x46/0xb0

    This happens because iocg->ioc is NULL. The field is initialized by
```

```
    ioc_pd_init() and never cleared. The NULL deref is caused by
    blkcg_activate_policy() installing blkg_policy_data before initializing it.

    blkcg_activate_policy() was doing the following:

    1. Allocate pd's for all existing blkg's and install them in blkg->pd[].
    2. Initialize all pd's.
    3. Online all pd's.

    blkcg_activate_policy() only grabs the queue_lock and may release and
    re-acquire the lock as allocation may need to sleep. ioc_weight_write()
    grabs blkcg->lock and iterates all its blkg's. The two can race and if
    ioc_weight_write() runs during #1 or between #1 and #2, it can encounter a
    pd which is not initialized yet, leading to crash.

    The crash can be reproduced with the following script:

      #!/bin/bash

      echo +io > /sys/fs/cgroup/cgroup.subtree_control
      systemd-run --unit touch-sda --scope dd if=/dev/sda of=/dev/null bs=1M count=1
iflag=direct
      echo 100 > /sys/fs/cgroup/system.slice/io.weight
      bash -c "echo '8:0 enable=1' > /sys/fs/cgroup/io.cost.qos" &
      sleep .2
      echo 100 > /sys/fs/cgroup/system.slice/io.weight

    with the following patch applied:

    > diff --git a/block/blk-cgroup.c b/block/blk-cgroup.c
    > index fc49be622e05..38d671d5e10c 100644
    > --- a/block/blk-cgroup.c
    > +++ b/block/blk-cgroup.c
    > @@ -1553,6 +1553,12 @@ int blkcg_activate_policy(struct gendisk *disk, const struct
blkcg_policy *pol)
    >               pd->online = false;
    >       }
    >
    > +       if (system_state == SYSTEM_RUNNING) {
    > +               spin_unlock_irq(&q->queue_lock);
    > +               ssleep(1);
    > +               spin_lock_irq(&q->queue_lock);
    > +       }
    > +
    >       /* all allocated, init in the same order */
    >       if (pol->pd_init_fn)
    >               list_for_each_entry_reverse(blkg, &q->blkg_list, q_node)

    I don't see a reason why all pd's should be allocated, initialized and
    onlined together. The only ordering requirement is that parent blkgs to be
    initialized and onlined before children, which is guaranteed from the
    walking order. Let's fix the bug by allocating, initializing and onlining pd
    for each blkg and holding blkcg->lock over initialization and onlining. This
    ensures that an installed blkg is always fully initialized and onlined
    removing the the race window.

    Signed-off-by: Tejun Heo <tj@kernel.org>
    Reported-by: Breno Leitao <leitao@debian.org>
    Fixes: 9d179b865449 ("blkcg: Fix multiple bugs in blkcg_activate_policy()")
    Link: https://lore.kernel.org/r/ZN0p5_W-Q9mAHBVY@slm.duckdns.org
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 02c126aa0b33ccf9e3c9b7cc5e436e3ccf564be9
Author: Vlad Karpovich <vkarpovi@opensource.cirrus.com>
```

```
Date:   Tue Aug 15 12:29:08 2023 -0500

    firmware: cs_dsp: Fix new control name check

    [ Upstream commit 7ac1102b227b36550452b663fd39ab1c09378a95 ]

    Before adding a new FW control, its name is checked against
    existing controls list. But the string length in strncmp used
    to compare controls names is taken from the list, so if beginnings
    of the controls are matching,  then the new control is not created.
    For example, if CAL_R control already exists, CAL_R_SELECTED
    is not created.
    The fix is to compare string lengths as well.

    Fixes: 6477960755fb ("ASoC: wm_adsp: Move check for control existence")
    Signed-off-by: Vlad Karpovich <vkarpovi@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20230815172908.3454056-1-vkarpovi@opensource.cirrus.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e46b2e7be8059d156af8c011dd8d665229b65886
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Tue Aug 8 18:49:12 2023 +0800

    md/raid5-cache: fix null-ptr-deref for r5l_flush_stripe_to_raid()

    [ Upstream commit 0d0bd28c500173bfca78aa840f8f36d261ef1765 ]

    r5l_flush_stripe_to_raid() will check if the list 'flushing_ios' is
    empty, and then submit 'flush_bio', however, r5l_log_flush_endio()
    is clearing the list first and then clear the bio, which will cause
    null-ptr-deref:

    T1: submit flush io
    raid5d
     handle_active_stripes
      r5l_flush_stripe_to_raid
       // list is empty
       // add 'io_end_ios' to the list
       bio_init
       submit_bio
       // io1

    T2: io1 is done
    r5l_log_flush_endio
     list_splice_tail_init
     // clear the list
                             T3: submit new flush io
                             ...
                             r5l_flush_stripe_to_raid
                              // list is empty
                              // add 'io_end_ios' to the list
                              bio_init
     bio_uninit
     // clear bio->bi_blkg
                              submit_bio
                              // null-ptr-deref

    Fix this problem by clearing bio before clearing the list in
    r5l_log_flush_endio().

    Fixes: 0dd00cba99c3 ("raid5-cache: fully initialize flush_bio when needed")
    Reported-and-tested-by: Corey Hickey <bugfood-ml@fatooh.org>
    Closes: https://lore.kernel.org/all/cddd7213-3dfd-4ab7-a3ac-edd54d74a626@fatooh.org/
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
```

```
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c406984738215dc20ac2dc63e49d70f20797730e
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Sat Jul 8 17:17:27 2023 +0800

    md/raid5-cache: fix a deadlock in r5l_exit_log()

    [ Upstream commit a705b11b358dee677aad80630e7608b2d5f56691 ]

    Commit b13015af94cf ("md/raid5-cache: Clear conf->log after finishing
    work") introduce a new problem:

    // caller hold reconfig_mutex
    r5l_exit_log
     flush_work(&log->disable_writeback_work)
                            r5c_disable_writeback_async
                             wait_event
                              /*
                               * conf->log is not NULL, and mddev_trylock()
                               * will fail, wait_event() can never pass.
                               */
     conf->log = NULL

    Fix this problem by setting 'config->log' to NULL before wake_up() as it
    used to be, so that wait_event() from r5c_disable_writeback_async() can
    exist. In the meantime, move forward md_unregister_thread() so that
    null-ptr-deref this commit fixed can still be fixed.

    Fixes: b13015af94cf ("md/raid5-cache: Clear conf->log after finishing work")
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230708091727.1417894-1-yukuai1@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 795fe88f0e8889a159657d856a7e6833af379e97
Author: Damien Le Moal <dlemoal@kernel.org>
Date:   Tue Aug 15 06:58:32 2023 +0900

    block: uapi: Fix compilation errors using ioprio.h with C++

    [ Upstream commit c7b4b23b36edf32239e7fc3b922797ff1d32b072 ]

    The use of the "class" argument name in the ioprio_value() inline
    function in include/uapi/linux/ioprio.h confuses C++ compilers
    resulting in compilation errors such as:

    /usr/include/linux/ioprio.h:110:43: error: expected primary-expression before ‘int’
      110 | static __always_inline __u16 ioprio_value(int class, int level, int hint)
          |                                           ^~~

    for user C++ programs including linux/ioprio.h.

    Avoid these errors by renaming the arguments of the ioprio_value()
    function to prioclass, priolevel and priohint. For consistency, the
    arguments of the IOPRIO_PRIO_VALUE() and IOPRIO_PRIO_VALUE_HINT() macros
    are also renamed in the same manner.

    Reported-by: Igor Pylypiv <ipylypiv@google.com>
    Fixes: 01584c1e2337 ("scsi: block: Improve ioprio value validity checks")
    Signed-off-by: Damien Le Moal <dlemoal@kernel.org>
    Reviewed-by: Chaitanya Kulkarni <kch@nvidia.com>
    Tested-by: Igor Pylypiv <ipylypiv@google.com>
```

```
    Link: https://lore.kernel.org/r/20230814215833.259286-1-dlemoal@kernel.org
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1458bf509dd01bd27517cb5cac6acffc5a94da1
Author: Tony Lindgren <tony@atomide.com>
Date:   Tue Aug 15 08:49:05 2023 +0300

    bus: ti-sysc: Fix cast to enum warning

    [ Upstream commit de44bf2f7683347f75690ef6cf61a1d5ba8f0891 ]

    Fix warning for "cast to smaller integer type 'enum sysc_soc' from 'const
    void *'".

    Cc: Nishanth Menon <nm@ti.com>
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202308150723.ziuGCdM3-lkp@intel.com/
    Fixes: e1e1e9bb9d94 ("bus: ti-sysc: Fix build warning for 64-bit build")
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cc5bd3035ec2ce50c48b04e4b020bd985fbfe802
Author: Rob Clark <robdclark@chromium.org>
Date:   Fri Aug 11 09:05:03 2023 -0700

    drm/msm/a690: Switch to a660_gmu.bin

    [ Upstream commit 18ff50e582a08eb365729b7c5507a86c41f2edf8 ]

    There isn't actually a a690_gmu.bin.  But it appears that the normal
    a660_gmu.bin works fine.  Normally all the devices within a sub-
    generation (or "family") will use the same fw, and a690 is in the a660
    family.

    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Fixes: 5e7665b5e484 ("drm/msm/adreno: Add Adreno A690 support")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/552406/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc23a39db8aae0274ac2d1773c3a81826bed7fab
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Aug 3 15:05:26 2023 +0200

    arm64: dts: qcom: sc8280xp-x13s: Unreserve NC pins

    [ Upstream commit 7868ed0144b33903e16a50485775f669c109e41a ]

    Pins 83-86 and 158-160 are NC, so there's no point in keeping them
    reserved. Take care of that.

    Fixes: 32c231385ed4 ("arm64: dts: qcom: sc8280xp: add Lenovo Thinkpad X13s devicetree")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230803-topic-x13s_pin-v1-1-fae792274e89@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 23b5f6686c6264c68b5746e01c0ebf2c8ce07a11
Author: David Wronek <davidwronek@gmail.com>
Date:   Sat Aug 5 15:09:37 2023 +0200

    arm64: dts: qcom: msm8996: Fix dsi1 interrupts
```

```
    [ Upstream commit bd3b4ac11845b428996cfd2c7b8302ba6a07340d ]

    Fix IRQ flags mismatch which was keeping dsi1 from probing by changing
    interrupts = <4> to interrupts = <5>.

    Fixes: 2752bb7d9b58 ("arm64: dts: qcom: msm8996: add second DSI interface")
    Signed-off-by: David Wronek <davidwronek@gmail.com>
    Acked-by: Yassine Oudjana <y.oudjana@protonmail.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230805130936.359860-2-davidwronek@gmail.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4714832a86787a8a714057cbaea199b45044076
Author: Rohit Agarwal <quic_rohiagar@quicinc.com>
Date:   Mon Aug 7 19:08:51 2023 +0530

    ARM: dts: qcom: sdx65-mtp: Update the pmic used in sdx65

    [ Upstream commit f636d6c356b339b0d29eed025f8bf9efcb6eb274 ]

    Update the pmic used in sdx65 platform to pm7250b.

    Fixes: 26380f298b2b (ARM: dts: qcom: sdx65-mtp: Add pmk8350b and pm8150b pmic)
    Signed-off-by: Rohit Agarwal <quic_rohiagar@quicinc.com>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/1691415534-31820-7-git-send-email-
quic_rohiagar@quicinc.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 11d0ca7002b34249e1d548eabce349412aa4cdbd
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Aug 9 21:20:25 2023 +0200

    arm64: dts: qcom: msm8998: Add missing power domain to MMSS SMMU

    [ Upstream commit 7f828f3207142351750e9545527341425187de7b ]

    The MMSS SMMU has its own power domain. Attach it so that we can drop
    the "keep it always-on" hack.

    Fixes: 05ce21b54423 ("arm64: dts: qcom: msm8998: Configure the multimedia subsystem
iommu")
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230531-topic-8998_mmssclk-v3-2-ba1b1fd9ee75@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8dcc78ddea52c48e85b400e63f238f36b3e4ae9a
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Aug 9 21:20:24 2023 +0200

    arm64: dts: qcom: msm8998: Drop bus clock reference from MMSS SMMU

    [ Upstream commit a3ce236364b82688ca4c7605f63c4efd68e9589c ]

    The MMSS SMMU has been abusingly consuming the exposed RPM interconnect
    clock. Drop it.

    Fixes: 05ce21b54423 ("arm64: dts: qcom: msm8998: Configure the multimedia subsystem
iommu")
    Reviewed-by: Jeffrey Hugo <quic_jhugo@quicinc.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
```

```
    Link: https://lore.kernel.org/r/20230531-topic-8998_mmssclk-v3-1-ba1b1fd9ee75@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3993a56cfc564bc881de589680d66a36413dcb41
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Sat Aug 12 00:47:36 2023 +0100

    arm64: dts: qcom: apq8016-sbc: Rename ov5640 enable-gpios to powerdown-gpios

    [ Upstream commit 4facccb44a82129195878750eed8f9890091c1b8 ]

    There are two control lines controlled by GPIO going into ov5640

    - Reset
    - Powerdown

    The driver and yaml expect "reset-gpios" and "powerdown-gpios" there has
    never been an "enable-gpios".

    Fixes: 39e0ce6cd1bf ("arm64: dts: qcom: apq8016-sbc: Add CCI/Sensor nodes")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230811234738.2859417-6-bryan.odonoghue@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b19598f32e3b48c4336b1585f4343edbe768fc0
Author: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
Date:   Sat Aug 12 00:47:33 2023 +0100

    arm64: dts: qcom: apq8016-sbc: Fix ov5640 regulator supply names

    [ Upstream commit 43a684580819e7f35b6cb38236be63c4cba26ef4 ]

    The ov5640 driver expects DOVDD, AVDD and DVDD as regulator supply names.

    The ov5640 has depended on these names since the driver was committed
    upstream in 2017. Similarly apq8016-sbc.dtsi has had completely different
    regulator names since its own initial commit in 2020.

    Perhaps the regulators were left on in previous 410c bootloaders. In any
    case today on 6.5 we won't switch on the ov5640 without correctly naming
    the regulators.

    Fixes: 39e0ce6cd1bf ("arm64: dts: qcom: apq8016-sbc: Add CCI/Sensor nodes")
    Signed-off-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230811234738.2859417-3-bryan.odonoghue@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ae0a8e23115c78a5ec3d22f97c5883bcfdbdc2ec
Author: Jason-JH.Lin <jason-jh.lin@mediatek.com>
Date:   Wed Jun 21 15:54:21 2023 +0800

    drm/mediatek: Fix void-pointer-to-enum-cast warning

    [ Upstream commit 89cba955f879b1c6a9a71f67c8fb92ea8f5dfdc4 ]

    1. Fix build warning message in mtk_disp_ovl_adaptor.c
    >> drivers/gpu/drm/mediatek/mtk_disp_ovl_adaptor.c:415:10:
      warning: cast to smaller integer type 'enum mtk_ovl_adaptor_comp_type'
      from 'const void *' [-Wvoid-pointer-to-enum-cast]
```

```
       type = (enum mtk_ovl_adaptor_comp_type)of_id->data;

              ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
              1 warning generated.

    2. Also fix the same warning message in mtk_drm_drv.c
    >> drivers/gpu/drm/mediatek/mtk_drm_drv.c:832:15:
       warning: cast to smaller integer type 'enum mtk_ddp_comp_type'
       from 'const void *' [-Wvoid-pointer-to-enum-cast]

       comp_type = (enum mtk_ddp_comp_type)of_id->data;

                   ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
                   1 warning generated.

    Signed-off-by: Jason-JH.Lin <jason-jh.lin@mediatek.com>
    Fixes: 453c3364632a ("drm/mediatek: Add ovl_adaptor support for MT8195")
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202305042054.ZtWME9OU-lkp@intel.com/
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230621075421.1982-1-jason-
jh.lin@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b9782d58dd5dab7cd7a78911691f81b32071bbb
Author: Sui Jingfeng <suijingfeng@loongson.cn>
Date:   Thu Jul 6 21:40:00 2023 +0800

    drm/mediatek: Fix potential memory leak if vmap() fail

    [ Upstream commit 379091e0f6d179d1a084c65de90fa44583b14a70 ]

    Also return -ENOMEM if such a failure happens, the implement should take
    responsibility for the error handling.

    Fixes: 3df64d7b0a4f ("drm/mediatek: Implement gem prime vmap/vunmap function")
    Reviewed-by: Matthias Brugger <matthias.bgg@gmail.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Signed-off-by: Sui Jingfeng <suijingfeng@loongson.cn>
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230706134000.130098-1-
suijingfeng@loongson.cn/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb836df627222eff5c711dc4d49c04c56ff12e80
Author: Aradhya Bhatia <a-bhatia1@ti.com>
Date:   Wed Aug 9 14:15:54 2023 +0530

    arm64: dts: ti: k3-am62x-sk-common: Update main-i2c1 frequency

    [ Upstream commit 73387da70f9c26b6fba4f62371d013cce14663d9 ]

    The Display Data Channel (DDC) transactions between an HDMI transmitter
    (SIL9022A in this case) and an HDMI monitor, occur at a maximum of
    100KHz. That's the maximum supported frequency within DDC standards.

    While the SIL9022A can transact with the core at 400KHz, it needs to
    drop the frequency to 100KHz when communicating with the monitor,
    otherwise, the i2c controller times out and shows warning like this.

    [  985.773431] omap_i2c 20010000.i2c: controller timed out
```

```
    That feature, however, has not been enabled in the SIL9022 driver.

    Since, dropping the frequency doesn't affect any other devices on the
    bus, drop the main-i2c1 frequency from 400KHz to 100KHz.

    Fixes: a841581451af ("arm64: dts: ti: Refractor AM625 SK dts")
    Signed-off-by: Aradhya Bhatia <a-bhatia1@ti.com>
    Link: https://lore.kernel.org/r/20230809084559.17322-2-a-bhatia1@ti.com
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d13b9d2612ba2e8bb40481b79c063c5f29250802
Author: Robert Marko <robert.marko@sartura.hr>
Date:   Fri Aug 11 13:01:16 2023 +0200

    ARM: dts: qcom: ipq4019: correct SDHCI XO clock

    [ Upstream commit b5ed7a5c1fdb3981713f7b637b72aa390c3db036 ]

    Using GCC_DCD_XO_CLK as the XO clock for SDHCI controller is not correct,
    it seems that I somehow made a mistake of passing it instead of the fixed
    XO clock.

    Fixes: 04b3b72b5b8f ("ARM: dts: qcom: ipq4019: Add SDHCI controller node")
    Signed-off-by: Robert Marko <robert.marko@sartura.hr>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230811110150.229966-1-robert.marko@sartura.hr
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 981b94165740cd735c6d670227f10b277cbefba1
Author: Marco Felsch <m.felsch@pengutronix.de>
Date:   Wed Aug 9 09:10:23 2023 +0200

    arm64: dts: imx8mp-debix: remove unused fec pinctrl node

    [ Upstream commit 574e4099d787c2eb41a43f14c453e422515bf658 ]

    The SoM A uses the EQOS ethernet interface and not the FEC, so drop the
    interface pinctrl node from the device tree.

    Fixes: c86d350aae68 ("arm64: dts: Add device tree for the Debix Model A Board")
    Signed-off-by: Marco Felsch <m.felsch@pengutronix.de>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Shawn Guo <shawnguo@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a933359cd800436a5b5c24aead1a3b15d35a56d
Author: Jason-JH.Lin <jason-jh.lin@mediatek.com>
Date:   Fri Jul 14 17:49:06 2023 +0800

    drm/mediatek: Add cnt checking for coverity issue

    [ Upstream commit d761b9450e31e5abd212f0085d424ed32760de5a ]

    CERT-C Characters and Strings (CERT STR31-C)
    all_drm_priv[cnt] evaluates to an address that could be at negative
    offset of an array.

    In mtk_drm_get_all_drm_priv():
    Guarantee that storage for strings has sufficient space for character
    data and the null terminator.

    So change cnt to unsigned int and check its max value.
```

```
    Fixes: 1ef7ed48356c ("drm/mediatek: Modify mediatek-drm for mt8195 multi mmsys support")
    Signed-off-by: Jason-JH.Lin <jason-jh.lin@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230714094908.13087-3-jason-
jh.lin@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d8de33248d364279cdf85d6a766def2350c07b65
Author: Jason-JH.Lin <jason-jh.lin@mediatek.com>
Date:   Fri Jul 14 17:49:05 2023 +0800

    drm/mediatek: Remove freeing not dynamic allocated memory

    [ Upstream commit 27b9e2ea3f2757da26bb8280e46f7fdbb1acb219 ]

    Fixing the coverity issue of:
    mtk_drm_cmdq_pkt_destroy frees address of mtk_crtc->cmdq_handle

    So remove the free function.

    Fixes: 7627122fd1c0 ("drm/mediatek: Add cmdq_handle in mtk_crtc")
    Signed-off-by: Jason-JH.Lin <jason-jh.lin@mediatek.com>
    Reviewed-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230714094908.13087-2-jason-
jh.lin@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9046b97555619ef56e917c958c847db8d414672a
Author: Tony Lindgren <tony@atomide.com>
Date:   Fri Aug 4 13:38:01 2023 +0300

    bus: ti-sysc: Fix build warning for 64-bit build

    [ Upstream commit e1e1e9bb9d943ec690670a609a5f660ca10eaf85 ]

    Fix "warning: cast from pointer to integer of different size" on 64-bit
    builds.

    Note that this is a cosmetic fix at this point as the driver is not yet
    used for 64-bit systems.

    Fixes: feaa8baee82a ("bus: ti-sysc: Implement SoC revision handling")
    Reviewed-by: Dhruva Gole <d-gole@ti.com>
    Reviewed-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Tony Lindgren <tony@atomide.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15795235e8f266c0deeda3ec565f0f326b87cbc2
Author: Apelete Seketeli <aseketeli@baylibre.com>
Date:   Thu Aug 10 01:16:22 2023 +0530

    arm64: dts: ti: k3-j784s4: Fix interrupt ranges for wkup & main gpio

    [ Upstream commit 05a1f130101e7a49ff1e8734939facd43596ea26 ]

    This patch fixes the interrupt range for wakeup and main domain gpio
    interrupt routers. They were wrongly subtracted by 32 instead of
    following what is defined in the interrupt map in the TRM (Table 9-35).
```

```
    Link:  http://www.ti.com/lit/pdf/spruj52
    Fixes: 4664ebd8346a ("arm64: dts: ti: Add initial support for J784S4 SoC")
    Signed-off-by: Apelete Seketeli <aseketeli@baylibre.com>
    Signed-off-by: Esteban Blanc <eblanc@baylibre.com>
    Signed-off-by: Jai Luthra <j-luthra@ti.com>
    Link: https://lore.kernel.org/r/20230810-tps6594-v6-4-2b2e2399e2ef@ti.com
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00855d02e0394cfdc99879ecd78fbcdd84ce06bc
Author: Udit Kumar <u-kumar1@ti.com>
Date:   Wed Aug 9 10:31:08 2023 +0530

    arm64: dts: ti: k3-j784s4-evm: Correct Pin mux offset for ADC

    [ Upstream commit 8be3ac2d8bd77bb9cb9ddbb7a545decf9f5e4181 ]

    After splitting wkup_pmx pin mux for J784S4 into four regions.
    Pin mux offset for ADC nodes were not updated to align with new
    regions, due to this while probing ADC driver out of range
    error was seen.

    Pin mux offsets for ADC nodes are corrected in this patch.

    Fixes: 14462bd0b247 ("arm64: dts: ti: k3-j784s4: Fix wakeup pinmux range and pinctrl
node offsets")
    Signed-off-by: Udit Kumar <u-kumar1@ti.com>
    Reviewed-by: Vaishnav Achath <vaishnav.a@ti.com>
    Link: https://lore.kernel.org/r/20230809050108.751164-1-u-kumar1@ti.com
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 26ec06978bf8ba09afa70c47813b863cee0ae013
Author: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
Date:   Tue Jul 25 09:32:24 2023 +0200

    drm/mediatek: dp: Add missing error checks in mtk_dp_parse_capabilities

    [ Upstream commit cfc146137a9f12e883ba64bc496b6da4d23f26d5 ]

    If reading the RX capabilities fails the training pattern will be set
    wrongly: add error checking for drm_dp_read_dpcd_caps() and return if
    anything went wrong with it.

    While at it, also add a less critical error check when writing to
    clear the ESI0 IRQ vector.

    Fixes: f70ac097a2cf ("drm/mediatek: Add MT8195 Embedded DisplayPort driver")
    Signed-off-by: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Tested-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: Alexandre Mergnat <amergnat@baylibre.com>
    Reviewed-by: CK Hu <ck.hu@mediatek.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230725073234.55892-2-
angelogioacchino.delregno@collabora.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 14a1b55482d8f0d6aaea5e6c1a0536d3ccfbbd2d
Author: Jinyoung Choi <j-young.choi@samsung.com>
Date:   Thu Aug 3 11:52:02 2023 +0900

    bio-integrity: create multi-page bvecs in bio_integrity_add_page()

    [ Upstream commit 0ece1d649b6dd615925a72bc1824d6b9fa5b998a ]
```

```
    In general, the bvec data structure consists of one for physically
    continuous pages. But, in the bvec configuration for bip, physically
    continuous integrity pages are composed of each bvec.

    Allow bio_integrity_add_page() to create multi-page bvecs, just like
    the bio payloads. This simplifies adding larger payloads, and fixes
    support for non-tiny workloads with nvme, which stopped using
    scatterlist for metadata a while ago.

    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Martin K. Petersen <martin.petersen@oracle.com>

    Fixes: 783b94bd9250 ("nvme-pci: do not build a scatterlist to map metadata")
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jinyoung Choi <j-young.choi@samsung.com>
    Tested-by: "Martin K. Petersen" <martin.petersen@oracle.com>
    Reviewed-by: "Martin K. Petersen" <martin.petersen@oracle.com>
    Link:
https://lore.kernel.org/r/20230803025202epcms2p82f57cbfe32195da38c776377b55aed59@epcms2p8
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 618bd48d0f0e3141bf498cbb8876f39b78a8b666
Author: Jinyoung Choi <j-young.choi@samsung.com>
Date:   Thu Aug 3 11:48:27 2023 +0900

    block: make bvec_try_merge_hw_page() non-static

    [ Upstream commit 7c8998f75d2d42ddefb172239b0f689392958309 ]

    This will be used for multi-page configuration for integrity payload.

    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Martin K. Petersen <martin.petersen@oracle.com>

    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Jinyoung Choi <j-young.choi@samsung.com>
    Tested-by: "Martin K. Petersen" <martin.petersen@oracle.com>
    Reviewed-by: "Martin K. Petersen" <martin.petersen@oracle.com>
    Link:
https://lore.kernel.org/r/20230803024827epcms2p838d9e9131492c86a159fff25d195658f@epcms2p8
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4a929e0a03308b9f16dc9003f98561deb7fbc081
Author: Christoph Hellwig <hch@lst.de>
Date:   Mon Jul 24 09:54:33 2023 -0700

    block: don't pass a bio to bio_try_merge_hw_seg

    [ Upstream commit ae42f0b3bf65912e122fc2e8d5f6d94b51156dba ]

    There is no good reason to pass the bio to bio_try_merge_hw_seg.  Just
    pass the current bvec and rename the function to bvec_try_merge_hw_page.
    This will allow reusing this function for supporting multi-page integrity
    payload bvecs.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Jinyoung Choi <j-young.choi@samsung.com>
    Link: https://lore.kernel.org/r/20230724165433.117645-9-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
```

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

commit 02ecc47f878172ef1d043b0fbee28dca22dac7dc
Author: Christoph Hellwig <hch@lst.de>
Date:   Mon Jul 24 09:54:32 2023 -0700

    block: move the bi_size update out of __bio_try_merge_page

    [ Upstream commit 858c708d9efb7e8e5c6320793b778cc17cf8368a ]

    The update of bi_size is the only thing in __bio_try_merge_page that
    needs a bio.  Move it to the callers, and merge __bio_try_merge_page
    and page_is_mergeable into a single bvec_try_merge_page that only takes
    the current bvec instead of a full bio.  This will allow reusing this
    function for supporting multi-page integrity payload bvecs.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Jinyoung Choi <j-young.choi@samsung.com>
    Link: https://lore.kernel.org/r/20230724165433.117645-8-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7f6e836590243209c450737ea0ccd634478e1527
Author: Christoph Hellwig <hch@lst.de>
Date:   Mon Jul 24 09:54:30 2023 -0700

    block: move the bi_size overflow check in __bio_try_merge_page

    [ Upstream commit 613699050a49760f1d70c74f71bd0b013ca3c356 ]

    Checking for availability in bi_size in a function that attempts to
    merge into an existing segment is a bit odd, as the limit also applies
    when adding a new segment.  This code works fine as we always call
    __bio_try_merge_page, but contributes to sub-optimal calling conventions
    and doesn't lead to clear code.

    Move it to two of the callers instead, the third one already has a more
    strict check that includes max_hw_segments anyway.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Jinyoung Choi <j-young.choi@samsung.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Link: https://lore.kernel.org/r/20230724165433.117645-6-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88257b923bceddbe678c90ef47d7425e3748255a
Author: Christoph Hellwig <hch@lst.de>
Date:   Mon Jul 24 09:54:29 2023 -0700

    block: move the bi_vcnt check out of __bio_try_merge_page

    [ Upstream commit 0eca8b6f97ac705c5806f7d062207379094fb114 ]

    Move the bi_vcnt out of __bio_try_merge_page and into the two callers
    that don't already have it in preparation for additional changes to
    __bio_try_merge_page.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Jinyoung Choi <j-young.choi@samsung.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
```

```
    Link: https://lore.kernel.org/r/20230724165433.117645-5-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b0c4c840b11d95c09bbe4ad7966f7e05d332e5d
Author: Christoph Hellwig <hch@lst.de>
Date:   Mon Jul 24 09:54:28 2023 -0700

    block: move the BIO_CLONED checks out of __bio_try_merge_page

    [ Upstream commit 939e1a370330841b2c0292a483d7b38f3ee45f88 ]

    __bio_try_merge_page is a way too low-level helper to assert that the
    bio is not cloned.  Move the check into bio_add_page and
    bio_iov_iter_get_pages instead, which are the high level entry points
    that should enforce this variant.  bio_add_hw_page already this
    check, coverig the third (indirect) caller of __bio_try_merge_page.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Reviewed-by: Jinyoung Choi <j-young.choi@samsung.com>
    Reviewed-by: Johannes Thumshirn <johannes.thumshirn@wdc.com>
    Link: https://lore.kernel.org/r/20230724165433.117645-4-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 0ece1d649b6d ("bio-integrity: create multi-page bvecs in
bio_integrity_add_page()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa149eddad6d9d69905b16ee5f3127d3a635db2a
Author: Pavel Begunkov <asml.silence@gmail.com>
Date:   Wed Aug 9 13:21:41 2023 +0100

    io_uring: fix drain stalls by invalid SQE

    [ Upstream commit cfdbaa3a291d6fd2cb4a1a70d74e63b4abc2f5ec ]

    cq_extra is protected by ->completion_lock, which io_get_sqe() misses.
    The bug is harmless as it doesn't happen in real life, requires invalid
    SQ index array and racing with submission, and only messes up the
    userspace, i.e. stall requests execution but will be cleaned up on
    ring destruction.

    Fixes: 15641e427070f ("io_uring: don't cache number of dropped SQEs")
    Signed-off-by: Pavel Begunkov <asml.silence@gmail.com>
    Link:
https://lore.kernel.org/r/66096d54651b1a60534bb2023f2947f09f50ef73.1691538547.git.asml.silen
ce@gmail.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7539d14deedbbb722b900491ba9380ec165bff20
Author: Zhiguo Niu <zhiguo.niu@unisoc.com>
Date:   Thu Aug 3 19:12:42 2023 +0800

    block/mq-deadline: use correct way to throttling write requests

    [ Upstream commit d47f9717e5cfd0dd8c0ba2ecfa47c38d140f1bb6 ]

    The original formula was inaccurate:
    dd->async_depth = max(1UL, 3 * q->nr_requests / 4);

    For write requests, when we assign a tags from sched_tags,
    data->shallow_depth will be passed to sbitmap_find_bit,
    see the following code:
```

```
    nr = sbitmap_find_bit_in_word(&sb->map[index],
                                  min_t (unsigned int,
                                  __map_depth(sb, index),
                                  depth),
                                  alloc_hint, wrap);

    The smaller of data->shallow_depth and __map_depth(sb, index)
    will be used as the maximum range when allocating bits.

    For a mmc device (one hw queue, deadline I/O scheduler):
    q->nr_requests = sched_tags = 128, so according to the previous
    calculation method, dd->async_depth = data->shallow_depth = 96,
    and the platform is 64bits with 8 cpus, sched_tags.bitmap_tags.sb.shift=5,
    sb.maps[]=32/32/32/32, 32 is smaller than 96, whether it is a read or
    a write I/O, tags can be allocated to the maximum range each time,
    which has not throttling effect.

    In addition, refer to the methods of bfg/kyber I/O scheduler,
    limit ratiois are calculated base on sched_tags.bitmap_tags.sb.shift.

    This patch can throttle write requests really.

    Fixes: 07757588e507 ("block/mq-deadline: Reserve 25% of scheduler tags for synchronous
requests")

    Signed-off-by: Zhiguo Niu <zhiguo.niu@unisoc.com>
    Reviewed-by: Bart Van Assche <bvanassche@acm.org>
    Link: https://lore.kernel.org/r/1691061162-22898-1-git-send-email-zhiguo.niu@unisoc.com
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8e76b944a7b9bddef190ffe2e29c9ae342ab91ed
Author: Gaosheng Cui <cuigaosheng1@huawei.com>
Date:   Tue Aug 8 20:14:35 2023 +0800

    audit: fix possible soft lockup in __audit_inode_child()

    [ Upstream commit b59bc6e37237e37eadf50cd5de369e913f524463 ]

    Tracefs or debugfs maybe cause hundreds to thousands of PATH records,
    too many PATH records maybe cause soft lockup.

    For example:
      1. CONFIG_KASAN=y && CONFIG_PREEMPTION=n
      2. auditctl -a exit,always -S open -k key
      3. sysctl -w kernel.watchdog_thresh=5
      4. mkdir /sys/kernel/debug/tracing/instances/test

    There may be a soft lockup as follows:
      watchdog: BUG: soft lockup - CPU#45 stuck for 7s! [mkdir:15498]
      Kernel panic - not syncing: softlockup: hung tasks
      Call trace:
       dump_backtrace+0x0/0x30c
       show_stack+0x20/0x30
       dump_stack+0x11c/0x174
       panic+0x27c/0x494
       watchdog_timer_fn+0x2bc/0x390
       __run_hrtimer+0x148/0x4fc
       __hrtimer_run_queues+0x154/0x210
       hrtimer_interrupt+0x2c4/0x760
       arch_timer_handler_phys+0x48/0x60
       handle_percpu_devid_irq+0xe0/0x340
       __handle_domain_irq+0xbc/0x130
       gic_handle_irq+0x78/0x460
```

```
        el1_irq+0xb8/0x140
        __audit_inode_child+0x240/0x7bc
        tracefs_create_file+0x1b8/0x2a0
        trace_create_file+0x18/0x50
        event_create_dir+0x204/0x30c
        __trace_add_new_event+0xac/0x100
        event_trace_add_tracer+0xa0/0x130
        trace_array_create_dir+0x60/0x140
        trace_array_create+0x1e0/0x370
        instance_mkdir+0x90/0xd0
        tracefs_syscall_mkdir+0x68/0xa0
        vfs_mkdir+0x21c/0x34c
        do_mkdirat+0x1b4/0x1d4
        __arm64_sys_mkdirat+0x4c/0x60
        el0_svc_common.constprop.0+0xa8/0x240
        do_el0_svc+0x8c/0xc0
        el0_svc+0x20/0x30
        el0_sync_handler+0xb0/0xb4
        el0_sync+0x160/0x180

    Therefore, we add cond_resched() to __audit_inode_child() to fix it.

    Fixes: 5195d8e217a7 ("audit: dynamically allocate audit_names when not enough space is
in the names array")
    Signed-off-by: Gaosheng Cui <cuigaosheng1@huawei.com>
    Signed-off-by: Paul Moore <paul@paul-moore.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 89a729f648b039bd5663f2dc5247596e5942bce5
Author: Nancy.Lin <nancy.lin@mediatek.com>
Date:   Thu Aug 3 17:48:43 2023 +0800

    drm/mediatek: Fix uninitialized symbol

    [ Upstream commit 63ee9438f2aeffb2d1b2df2599c168ca08d35025 ]

    Fix Smatch static checker warning
      -Fix uninitialized symbol comp_pdev in mtk_ddp_comp_init.

    Fixes: 0d9eee9118b7 ("drm/mediatek: Add drm ovl_adaptor sub driver for MT8195")
    Signed-off-by: Nancy.Lin <nancy.lin@mediatek.com>
    Link: https://patchwork.kernel.org/project/dri-devel/patch/20230803094843.4439-1-
nancy.lin@mediatek.com/
    Signed-off-by: Chun-Kuang Hu <chunkuang.hu@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cf7b1dd2b38b14aee0b03388935d4eddd4d91c47
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Aug 7 16:09:41 2023 -0500

    ASoC: SOF: Intel: fix u16/32 confusion in LSDIID

    [ Upstream commit 7a52d7062e02af4a479da24b40cfd76b54c0cd6c ]

    Likely a combination of copy-paste and test coverage problem. Oops.

    Fixes: 87a6ddc0cf1c ("ASoC: SOF: Intel: hda-mlink: program SoundWire LSDIID registers")
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Link: https://lore.kernel.org/r/20230807210959.506849-3-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 59bf56fda7eb86570828833b9e385d0209c20e14
Author: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
Date:   Mon Aug 7 16:09:40 2023 -0500

    ASoC: SOF: Intel: hda-mlink: fix off-by-one error

    [ Upstream commit 7075b0c91b3cd5d32b4ac7403f771a3253d3fbf6 ]

    The HCHAN parameter should be the highest channel number, not the
    channel count.

    While we're at it, handle LCHAN with the dual __ffs helper.

    Fixes: ccc2f0c1b6b6 ("ASoC: SOF: Intel: hda-mlink: add helper to program SoundWire
PCMSyCM registers")
    Signed-off-by: Pierre-Louis Bossart <pierre-louis.bossart@linux.intel.com>
    Reviewed-by: Bard Liao <yung-chuan.liao@linux.intel.com>
    Reviewed-by: Rander Wang <rander.wang@intel.com>
    Link: https://lore.kernel.org/r/20230807210959.506849-2-pierre-
louis.bossart@linux.intel.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5a94f94b5622b5da68d5d0eda41a2f68a97505c8
Author: Rob Clark <robdclark@chromium.org>
Date:   Thu Aug 3 10:34:28 2023 -0700

    drm/msm/a6xx: Fix GMU lockdep splat

    [ Upstream commit 3136a0f83519076edfbc14be65f286785434189a ]

    For normal GPU devfreq, we need to acquire the GMU lock while already
    holding devfreq locks.  But in the teardown path, we were calling
    dev_pm_domain_detach() while already holding the GMU lock, resulting in
    this lockdep splat:

       ======================================================
       WARNING: possible circular locking dependency detected
       6.4.3-debug+ #3 Not tainted
       ------------------------------------------------------
       ring0/391 is trying to acquire lock:
       ffffff80a025c078 (&devfreq->lock){+.+.}-{3:3}, at: qos_notifier_call+0x30/0x74

       but task is already holding lock:
       ffffff809b8c1ce8 (&(c->notifiers)->rwsem){++++}-{3:3}, at:
blocking_notifier_call_chain+0x34/0x78

       which lock already depends on the new lock.

       the existing dependency chain (in reverse order) is:

       -> #4 (&(c->notifiers)->rwsem){++++}-{3:3}:
              down_write+0x58/0x74
              __blocking_notifier_chain_register+0x64/0x84
              blocking_notifier_chain_register+0x1c/0x28
              freq_qos_add_notifier+0x5c/0x7c
              dev_pm_qos_add_notifier+0xd4/0xf0
              devfreq_add_device+0x42c/0x560
              devm_devfreq_add_device+0x6c/0xb8
              msm_devfreq_init+0xa8/0x16c [msm]
              msm_gpu_init+0x368/0x54c [msm]
              adreno_gpu_init+0x248/0x2b0 [msm]
              a6xx_gpu_init+0x2d0/0x384 [msm]
              adreno_bind+0x264/0x2bc [msm]
              component_bind_all+0x124/0x1f4
```

```
               msm_drm_bind+0x2d0/0x5f4 [msm]
               try_to_bring_up_aggregate_device+0x88/0x1a4
               __component_add+0xd4/0x128
               component_add+0x1c/0x28
               dp_display_probe+0x37c/0x3c0 [msm]
               platform_probe+0x70/0xc0
               really_probe+0x148/0x280
               __driver_probe_device+0xfc/0x114
               driver_probe_device+0x44/0x100
               __device_attach_driver+0x64/0xdc
               bus_for_each_drv+0xb0/0xd8
               __device_attach+0xe4/0x140
               device_initial_probe+0x1c/0x28
               bus_probe_device+0x44/0xb0
               deferred_probe_work_func+0xb0/0xc8
               process_one_work+0x288/0x3d8
               worker_thread+0x1f0/0x260
               kthread+0xf0/0x100
               ret_from_fork+0x10/0x20

       -> #3 (dev_pm_qos_mtx){+.+.}-{3:3}:
               __mutex_lock+0xc8/0x388
               mutex_lock_nested+0x2c/0x38
               dev_pm_qos_remove_notifier+0x3c/0xc8
               genpd_remove_device+0x40/0x11c
               genpd_dev_pm_detach+0x88/0x130
               dev_pm_domain_detach+0x2c/0x3c
               a6xx_gmu_remove+0x44/0xdc [msm]
               a6xx_destroy+0x7c/0xa4 [msm]
               adreno_unbind+0x50/0x64 [msm]
               component_unbind+0x44/0x64
               component_unbind_all+0xb4/0xbc
               msm_drm_uninit.isra.0+0x124/0x17c [msm]
               msm_drm_bind+0x340/0x5f4 [msm]
               try_to_bring_up_aggregate_device+0x88/0x1a4
               __component_add+0xd4/0x128
               component_add+0x1c/0x28
               dp_display_probe+0x37c/0x3c0 [msm]
               platform_probe+0x70/0xc0
               really_probe+0x148/0x280
               __driver_probe_device+0xfc/0x114
               driver_probe_device+0x44/0x100
               __device_attach_driver+0x64/0xdc
               bus_for_each_drv+0xb0/0xd8
               __device_attach+0xe4/0x140
               device_initial_probe+0x1c/0x28
               bus_probe_device+0x44/0xb0
               deferred_probe_work_func+0xb0/0xc8
               process_one_work+0x288/0x3d8
               worker_thread+0x1f0/0x260
               kthread+0xf0/0x100
               ret_from_fork+0x10/0x20

       -> #2 (&a6xx_gpu->gmu.lock){+.+.}-{3:3}:
               __mutex_lock+0xc8/0x388
               mutex_lock_nested+0x2c/0x38
               a6xx_gpu_set_freq+0x38/0x64 [msm]
               msm_devfreq_target+0x170/0x18c [msm]
               devfreq_set_target+0x90/0x1e4
               devfreq_update_target+0xb4/0xf0
               update_devfreq+0x1c/0x28
               devfreq_monitor+0x3c/0x10c
               process_one_work+0x288/0x3d8
               worker_thread+0x1f0/0x260
               kthread+0xf0/0x100
```

```
               ret_from_fork+0x10/0x20

        -> #1 (&df->lock){+.+.}-{3:3}:
               __mutex_lock+0xc8/0x388
               mutex_lock_nested+0x2c/0x38
               msm_devfreq_get_dev_status+0x4c/0x104 [msm]
               devfreq_simple_ondemand_func+0x5c/0x128
               devfreq_update_target+0x68/0xf0
               update_devfreq+0x1c/0x28
               devfreq_monitor+0x3c/0x10c
               process_one_work+0x288/0x3d8
               worker_thread+0x1f0/0x260
               kthread+0xf0/0x100
               ret_from_fork+0x10/0x20

        -> #0 (&devfreq->lock){+.+.}-{3:3}:
               __lock_acquire+0xdf8/0x109c
               lock_acquire+0x234/0x284
               __mutex_lock+0xc8/0x388
               mutex_lock_nested+0x2c/0x38
               qos_notifier_call+0x30/0x74
               qos_min_notifier_call+0x1c/0x28
               notifier_call_chain+0xf4/0x114
               blocking_notifier_call_chain+0x4c/0x78
               pm_qos_update_target+0x184/0x190
               freq_qos_apply+0x4c/0x64
               apply_constraint+0xf8/0xfc
               __dev_pm_qos_update_request+0x138/0x164
               dev_pm_qos_update_request+0x44/0x68
               msm_devfreq_boost+0x40/0x70 [msm]
               msm_devfreq_active+0xc0/0xf0 [msm]
               msm_gpu_submit+0xc8/0x12c [msm]
               msm_job_run+0x88/0x128 [msm]
               drm_sched_main+0x240/0x324 [gpu_sched]
               kthread+0xf0/0x100
               ret_from_fork+0x10/0x20

        other info that might help us debug this:
        Chain exists of:
          &devfreq->lock --> dev_pm_qos_mtx --> &(c->notifiers)->rwsem
         Possible unsafe locking scenario:
               CPU0                    CPU1
               ----                    ----
          rlock(&(c->notifiers)->rwsem);
                                       lock(dev_pm_qos_mtx);
                                       lock(&(c->notifiers)->rwsem);
          lock(&devfreq->lock);

         *** DEADLOCK ***
        4 locks held by ring0/391:
         #0: ffffff809c811170 (&gpu->lock){+.+.}-{3:3}, at: msm_job_run+0x7c/0x128 [msm]
         #1: ffffff809c811208 (&gpu->active_lock){+.+.}-{3:3}, at: msm_gpu_submit+0xa8/0x12c
 [msm]
         #2: fffffecbbb46600 (dev_pm_qos_mtx){+.+.}-{3:3}, at:
 dev_pm_qos_update_request+0x38/0x68
         #3: ffffff809b8c1ce8 (&(c->notifiers)->rwsem){++++}-{3:3}, at:
 blocking_notifier_call_chain+0x34/0x78

        stack backtrace:
        CPU: 6 PID: 391 Comm: ring0 Not tainted 6.4.3debug+ #3
        Hardware name: Google Villager (rev1+) with LTE (DT)
        Call trace:
         dump_backtrace+0xb4/0xf0
         show_stack+0x20/0x30
         dump_stack_lvl+0x60/0x84
```

```
         dump_stack+0x18/0x24
         print_circular_bug+0x1cc/0x234
         check_noncircular+0x78/0xac
         __lock_acquire+0xdf8/0x109c
         lock_acquire+0x234/0x284
         __mutex_lock+0xc8/0x388
         mutex_lock_nested+0x2c/0x38
         qos_notifier_call+0x30/0x74
         qos_min_notifier_call+0x1c/0x28
         notifier_call_chain+0xf4/0x114
         blocking_notifier_call_chain+0x4c/0x78
         pm_qos_update_target+0x184/0x190
         freq_qos_apply+0x4c/0x64
         apply_constraint+0xf8/0xfc
         __dev_pm_qos_update_request+0x138/0x164
         dev_pm_qos_update_request+0x44/0x68
         msm_devfreq_boost+0x40/0x70 [msm]
         msm_devfreq_active+0xc0/0xf0 [msm]
         msm_gpu_submit+0xc8/0x12c [msm]
         msm_job_run+0x88/0x128 [msm]
         drm_sched_main+0x240/0x324 [gpu_sched]
         kthread+0xf0/0x100
         ret_from_fork+0x10/0x20

    Fix this by only synchronizing access to gmu->initialized.

    Fixes: 4cd15a3e8b36 ("drm/msm/a6xx: Make GPU destroy a bit safer")
    Cc: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Reviewed-by: Douglas Anderson <dianders@chromium.org>
    Patchwork: https://patchwork.freedesktop.org/patch/551171/
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 13868cae758da8907f2e6e32656943de2cbceef6
Author: Fabio Estevam <festevam@denx.de>
Date:   Tue Jun 20 20:23:19 2023 -0300

    drm/msm/a2xx: Call adreno_gpu_init() earlier

    [ Upstream commit db07ce5da8b26bfeaf437a676ae49bd3bb1eace6 ]

    The adreno_is_a20x() and adreno_is_a225() functions rely on the
    GPU revision, but such information is retrieved inside adreno_gpu_init(),
    which is called afterwards.

    Fix this problem by caling adreno_gpu_init() earlier, so that
    the GPU information revision is available when adreno_is_a20x()
    and adreno_is_a225() run.

    Tested on a imx53-qsb board.

    Fixes: 21af872cd8c6 ("drm/msm/adreno: add a2xx")
    Signed-off-by: Fabio Estevam <festevam@denx.de>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/543456/
    Signed-off-by: Rob Clark <robdclark@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f548c10b6e894064dcaf5410602f217e158385a
Author: Yang Wang <kevinyang.wang@amd.com>
Date:   Tue Aug 1 16:53:23 2023 +0800

    drm/amd/pm: fix variable dereferenced issue in amdgpu_device_attr_create()

    [ Upstream commit 25e6373a5b8efc623443f2699d2b929bf3067d76 ]
```

```
    - fix variable ('attr') dereferenced issue.
    - using condition check instead of BUG_ON().

    Fixes: 4e01847c38f7 ("drm/amdgpu: optimize amdgpu device attribute code")
    Cc: Dan Carpenter <dan.carpenter@linaro.org>
    Signed-off-by: Yang Wang <kevinyang.wang@amd.com>
    Reviewed-by: Kenneth Feng <kenneth.feng@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4885624e4a118e75c2cccf2359c4a620b48c4b5b
Author: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
Date:   Fri Jul 28 09:51:13 2023 +0530

    drm/amdgpu: Move vram, gtt & flash defines to amdgpu_ ttm & _psp.h

    [ Upstream commit 4e2abc197f11e25b5813d4c42dada19d36b04666 ]

    As amdgpu.h is getting decomposed, move vram and gtt extern defines into
    amdgpu_ttm.h & flash extern to amdgpu_psp.h

    Fixes: f9acfafc3458 ("drm/amdgpu: Move externs to amdgpu.h file from amdgpu_drv.c")
    Suggested-by: Christian König <christian.koenig@amd.com>
    Cc: Mario Limonciello <mario.limonciello@amd.com>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
    Acked-by: Guchun Chen <guchun.chen@amd.com>
    Reviewed-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 72c33625316e62d055f5d2198cdf4cba4df8d794
Author: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
Date:   Fri Jul 28 09:43:11 2023 +0530

    drm/amdgpu: Sort the includes in amdgpu/amdgpu_drv.c

    [ Upstream commit e2e42edfe8533af7b30f505d41d44e0d180065da ]

    Sort the include files that are included in amdgpu_drv.c alphabetically.

    Suggested-by: Mario Limonciello <mario.limonciello@amd.com>
    Cc: Mario Limonciello <mario.limonciello@amd.com>
    Cc: Christian König <christian.koenig@amd.com>
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
    Acked-by: Guchun Chen <guchun.chen@amd.com>
    Reviewed-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Stable-dep-of: 4e2abc197f11 ("drm/amdgpu: Move vram, gtt & flash defines to amdgpu_ ttm
& _psp.h")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 254fe9482d3a2fbc2e0d3da3ab5fe4b99abc29a3
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Thu Jul 6 08:52:39 2023 +0300

    smackfs: Prevent underflow in smk_set_cipso()

    [ Upstream commit 3ad49d37cf5759c3b8b68d02e3563f633d9c1aee ]

    There is a upper bound to "catlen" but no lower bound to prevent
    negatives.  I don't see that this necessarily causes a problem but we
```

```
    may as well be safe.

    Fixes: e114e473771c ("Smack: Simplified Mandatory Access Control Kernel")
    Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
    Signed-off-by: Casey Schaufler <casey@schaufler-ca.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 529a00c893cfc2751050032320c4a332437bf214
Author: Udit Kumar <u-kumar1@ti.com>
Date:   Fri Aug 4 13:23:41 2023 +0530

    arm64: dts: ti: k3-j721s2: correct pinmux offset for ospi

    [ Upstream commit 06c4e7aa4af0682910ea52d7c23d85f59ea7dcc6 ]

    Due to non-addressable regions in J721S2 SOC wkup_pmx was split
    into four regions from wkup_pmx0 to wkup_pmx3.

    Correcting OSPI1 pin mux, which now falls under wkup_pmx1.
    Along with that removing unused pin mux for OSPI-0.

    Fixes: 6bc829ceea41 ("arm64: dts: ti: k3-j721s2: Fix wkup pinmux range")

    Signed-off-by: Udit Kumar <u-kumar1@ti.com>
    Reviewed-by: Vaishnav Achath <vaishnav.a@ti.com>
    Link: https://lore.kernel.org/r/20230804075341.3858488-1-u-kumar1@ti.com
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a266e42ecc5e79ab1a84b2b117d2932b5af73ee5
Author: Udit Kumar <u-kumar1@ti.com>
Date:   Wed Aug 2 17:11:26 2023 +0530

    arm64: dts: ti: k3-j784s4-evm: Correct Pin mux offset for ospi

    [ Upstream commit f10f836ccfea21ae3ad3066eb9576625f1acdea2 ]

    After splitting wkup_pmx pin mux for J784S4 into four regions.
    Pin mux offset for OSPI nodes were not updated to align with new
    regions, due to this while setting ospi pin muxes out of range
    error was seen.

    Pin mux offsets for OSPI nodes are corrected in this patch.

    Fixes: 14462bd0b247 ("arm64: dts: ti: k3-j784s4: Fix wakeup pinmux range and pinctrl
node offsets")
    Signed-off-by: Udit Kumar <u-kumar1@ti.com>
    Tested-by: Vaishnav Achath <vaishnav.a@ti.com>
    Link: https://lore.kernel.org/r/20230802114126.162445-1-u-kumar1@ti.com
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 707b787b90668d79566b77362b162f9093531b3d
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Aug 2 13:04:19 2023 +0300

    drm/msm/dpu: fix the irq index in dpu_encoder_phys_wb_wait_for_commit_done

    [ Upstream commit d93cf453f51da168f4410ba73656f1e862096973 ]

    Since commit 1e7ac595fa46 ("drm/msm/dpu: pass irq to
    dpu_encoder_helper_wait_for_irq()") the
    dpu_encoder_phys_wb_wait_for_commit_done expects the IRQ index rather
    than the IRQ index in phys_enc->intr table, however writeback got the
    older invocation in place. This was unnoticed for several releases, but
```

```
    now it's time to fix it.

    Fixes: d7d0e73f7de3 ("drm/msm/dpu: introduce the dpu_encoder_phys_* for writeback")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/550924/
    Link: https://lore.kernel.org/r/20230802100426.4184892-2-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d6c4a1a4e6678cb98dd57964f133a995ecc91c1
Author: Zhang Shurong <zhang_shurong@foxmail.com>
Date:   Sat Jul 15 22:13:38 2023 +0800

    firmware: meson_sm: fix to avoid potential NULL pointer dereference

    [ Upstream commit f2ed165619c16577c02b703a114a1f6b52026df4 ]

    of_match_device() may fail and returns a NULL pointer.

    Fix this by checking the return value of of_match_device.

    Fixes: 8cde3c2153e8 ("firmware: meson_sm: Rework driver as a proper platform driver")
    Signed-off-by: Zhang Shurong <zhang_shurong@foxmail.com>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/tencent_AA08AAA6C4F34D53ADCE962E188A879B8206@qq.com
    Signed-off-by: Neil Armstrong <neil.armstrong@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 12dfd02cbd1a678fbd66be0c2f79d5299c4921a9
Author: Daniel Vetter <daniel.vetter@ffwll.ch>
Date:   Thu Aug 3 22:45:21 2023 +0200

    drm/msm/mdp5: Don't leak some plane state

    [ Upstream commit fd0ad3b2365c1c58aa5a761c18efc4817193beb6 ]

    Apparently no one noticed that mdp5 plane states leak like a sieve
    ever since we introduced plane_state->commit refcount a few years ago
    in 21a01abbe32a ("drm/atomic: Fix freeing connector/plane state too
    early by tracking commits, v3.")

    Fix it by using the right helpers.

    Fixes: 21a01abbe32a ("drm/atomic: Fix freeing connector/plane state too early by
tracking commits, v3.")
    Cc: Maarten Lankhorst <maarten.lankhorst@linux.intel.com>
    Cc: Daniel Vetter <daniel.vetter@ffwll.ch>
    Cc: Rob Clark <robdclark@gmail.com>
    Cc: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Cc: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Cc: linux-arm-msm@vger.kernel.org
    Cc: freedreno@lists.freedesktop.org
    Reported-and-tested-by: dorum@noisolation.com
    Cc: dorum@noisolation.com
    Signed-off-by: Daniel Vetter <daniel.vetter@intel.com>
    Reviewed-by: Rob Clark <robdclark@gmail.com>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/551236/
    Link: https://lore.kernel.org/r/20230803204521.928582-1-daniel.vetter@ffwll.ch
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ad14ed7403a9563ba6aef0baa425764b03ac047
Author: Chen Jiahao <chenjiahao16@huawei.com>
Date:   Tue Aug 1 17:48:07 2023 +0800
```

```
    soc: qcom: smem: Fix incompatible types in comparison

    [ Upstream commit 5f908786cf44fcb397cfe0f322ef2f41b0909e2a ]

    This patch fixes the following sparse error:

    drivers/soc/qcom/smem.c:738:30: error: incompatible types in comparison expression
 (different add        ress spaces):
    drivers/soc/qcom/smem.c:738:30:    void *
    drivers/soc/qcom/smem.c:738:30:    void [noderef] __iomem *

    In addr_in_range(), "base" is of type void __iomem *, converting
    void *addr to the same type to fix above sparse error.

    Fixes: 20bb6c9de1b7 ("soc: qcom: smem: map only partitions used by local HOST")
    Signed-off-by: Chen Jiahao <chenjiahao16@huawei.com>
    Link: https://lore.kernel.org/r/20230801094807.4146779-1-chenjiahao16@huawei.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c48b523b63571bd77644b8a30547c8147f6ddd35
Author: Abel Vesa <abel.vesa@linaro.org>
Date:   Tue Aug 1 12:52:46 2023 +0300

    arm64: dts: qcom: sm8550-mtp: Add missing supply for L1B regulator

    [ Upstream commit 2c9e45dfeed126488aa73e7b82b3576c4c6f1036 ]

    Even though currently there is no consumer for L1B, add the supply
    for it anyway.

    Fixes: 71342fb91eae ("arm64: dts: qcom: Add base SM8550 MTP dts")
    Signed-off-by: Abel Vesa <abel.vesa@linaro.org>
    Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
    Link: https://lore.kernel.org/r/20230801095246.2884770-1-abel.vesa@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee5c3905350666900da89e39ed8e7c4dec686e8b
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Aug 2 21:36:55 2023 +0300

    drm/msm/dpu: fix DSC 1.2 enc subblock length

    [ Upstream commit 57a1ca6cf73b164ff93c2a541a6fc2337fd07b20 ]

    Both struct dpu_dsc_sub_blks instances declare enc subblock length to be
    0x100, while the actual length is 0x9c (last register having offset 0x98).
    Reduce subblock length to remove the empty register space from being
    dumped.

    Fixes: 0d1b10c63346 ("drm/msm/dpu: add DSC 1.2 hw blocks for relevant chipsets")
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/550999/
    Link: https://lore.kernel.org/r/20230802183655.4188640-2-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f60f99031a32f883e4c597436561b23561424a55
Author: Ryan McCann <quic_rmccann@quicinc.com>
Date:   Fri Jul 7 18:24:42 2023 -0700

    drm/msm/dpu: Define names for unnamed sblks
```

```
    [ Upstream commit 46998bf8431c30879be29bbbf67eefb583136ccb ]

    Some sub-blocks in the hw catalog have not been given a name, so when the
    registers from that block are dumped, there is no name to reference.
    Define names for relevant sub-blocks to fix this.

    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Ryan McCann <quic_rmccann@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/546199/
    Link: https://lore.kernel.org/r/20230622-devcoredump_patch-v5-3-67e8b66c4723@quicinc.com
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Stable-dep-of: 57a1ca6cf73b ("drm/msm/dpu: fix DSC 1.2 enc subblock length")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 80b455f6f42fc5c117459bdd643e20b8203158f6
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Apr 20 03:14:54 2023 +0200

    drm/msm/dpu1: Rename sm8150_dspp_blk to sdm845_dspp_blk

    [ Upstream commit 9891b3df2b4300d24735c1a1822985d2d173aade ]

    SDM845 was the first SoC to include both PCC v4 and GC v1.8.
    We don't currently support any other blocks but the common config
    for these two can be reused for a large amount of SoCs.

    Rename it to indicate the origin of that combo.

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/533003/
    Link: https://lore.kernel.org/r/20230420-topic-dpu_gc-v1-1-d9d1a5e40917@linaro.org
    [DB: also applied to new catalog files]
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Stable-dep-of: 57a1ca6cf73b ("drm/msm/dpu: fix DSC 1.2 enc subblock length")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5c8f3f36cf6514fc1037816bdac2c584c63bd798
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Wed Aug 2 21:36:54 2023 +0300

    drm/msm/dpu: fix DSC 1.2 block lengths

    [ Upstream commit e550ad0e5c3d1a521413e6efb22729698a70110b ]

    All DSC_BLK_1_2 declarations incorrectly pass 0x29c as the block length.
    This includes the common block itself, enc subblocks and some empty
    space around. Change that to pass 0x4 instead, the length of common
    register block itself.

    Fixes: 0d1b10c63346 ("drm/msm/dpu: add DSC 1.2 hw blocks for relevant chipsets")
    Reported-by: Ryan McCann <quic_rmccann@quicinc.com>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Patchwork: https://patchwork.freedesktop.org/patch/550998/
    Link: https://lore.kernel.org/r/20230802183655.4188640-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e5090d00883f78d26baa5f9abeac794307b6324
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Tue Jul 4 05:21:30 2023 +0300
```

```
    drm/msm/dpu: inline DSC_BLK and DSC_BLK_1_2 macros

    [ Upstream commit 194347df5844e76bfc437e4aff2e1ece62af39c2 ]

    To simplify making changes to the hardware block definitions, expand
    corresponding macros. This way making all the changes are more obvious
    and visible in the source files.

    Tested-by: Marijn Suijten <marijn.suijten@somainline.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Patchwork: https://patchwork.freedesktop.org/patch/545370/
    Link: https://lore.kernel.org/r/20230704022136.130522-14-dmitry.baryshkov@linaro.org
    Stable-dep-of: e550ad0e5c3d ("drm/msm/dpu: fix DSC 1.2 block lengths")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f39562f04b0584c71f7cedd51158a3bf5a369f55
Author: Jonathan Marek <jonathan@marek.ca>
Date:   Wed Aug 2 09:48:53 2023 -0400

    drm/msm/dpu: increase memtype count to 16 for sm8550

    [ Upstream commit 42d0d253ed03b961c325ff756eec0480cb4adc6b ]

    sm8550 has 16 vbif clients.

    This fixes the extra 2 clients (DMA4/DMA5) not having their memtype
    initialized. This fixes DMA4/DMA5 planes not displaying correctly.

    Fixes: efcd0107727c ("drm/msm/dpu: add support for SM8550")
    Signed-off-by: Jonathan Marek <jonathan@marek.ca>
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Tested-by: Neil Armstrong <neil.armstrong@linaro.org> # on SM8550-QRD
    Patchwork: https://patchwork.freedesktop.org/patch/550968/
    Link: https://lore.kernel.org/r/20230802134900.30435-1-jonathan@marek.ca
    [DB: fixed the Fixes tag]
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 57b74bc86c5a5201b203b16c36fd3e585b071a69
Author: Jiasheng Jiang <jiasheng@iscas.ac.cn>
Date:   Wed Jun 7 10:05:29 2023 +0800

    drm: xlnx: zynqmp_dpsub: Add missing check for dma_set_mask

    [ Upstream commit 1832fba7f9780aff67c96ad30f397c2d76141833 ]

    Add check for dma_set_mask() and return the error if it fails.

    Fixes: d76271d22694 ("drm: xlnx: DRM/KMS driver for Xilinx ZynqMP DisplayPort
Subsystem")
    Signed-off-by: Jiasheng Jiang <jiasheng@iscas.ac.cn>
    Reviewed-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Reviewed-by: Tomi Valkeinen <tomi.valkeinen@ideasonboard.com>
    Signed-off-by: Laurent Pinchart <laurent.pinchart@ideasonboard.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a0b247ba37a2a27a4a251eb45114920b0bbb534b
Author: Nayna Jain <nayna@linux.ibm.com>
Date:   Tue Jul 11 12:44:47 2023 -0400

    ima: Remove deprecated IMA_TRUSTED_KEYRING Kconfig

    [ Upstream commit 5087fd9e80e539d2163accd045b73da64de7de95 ]

    Time to remove "IMA_TRUSTED_KEYRING".
```

```
    Fixes: f4dc37785e9b ("integrity: define '.evm' as a builtin 'trusted' keyring") # v4.5+
    Signed-off-by: Nayna Jain <nayna@linux.ibm.com>
    Signed-off-by: Mimi Zohar <zohar@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f3123adfa19fb8f7adea3bce6d6afaa5fcfb02b
Author: Marek Vasut <marex@denx.de>
Date:   Sun Jul 9 15:49:14 2023 +0200

    drm/panel: simple: Add missing connector type and pixel format for AUO T215HVN01

    [ Upstream commit 7a675a8fa598edb29a664a91adb80f0340649f6f ]

    The connector type and pixel format are missing for this panel,
    add them to prevent various drivers from failing to determine
    either of those parameters.

    Fixes: 7ee933a1d5c4 ("drm/panel: simple: Add support for AUO T215HVN01")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230709134914.449328-1-
marex@denx.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7faed9ac3a1c31404ae04bed448144c4a03aaa1
Author: Geert Uytterhoeven <geert@linux-m68k.org>
Date:   Thu Mar 17 09:18:30 2022 +0100

    drm/repaper: Reduce temporary buffer size in repaper_fb_dirty()

    [ Upstream commit fedf429e071f6dbbe7a69dfc342492e037692018 ]

    As the temporary buffer is no longer used to store 8-bit grayscale data,
    its size can be reduced to the size needed to store the monochrome
    bitmap data.

    Fixes: 24c6bedefbe71de9 ("drm/repaper: Use format helper for xrgb8888 to monochrome
conversion")
    Signed-off-by: Geert Uytterhoeven <geert@linux-m68k.org>
    Reviewed-by: Javier Martinez Canillas <javierm@redhat.com>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20220317081830.1211400-6-
geert@linux-m68k.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1cdb0e02fd4a3d8d7a7de58279b8b77cf4900f5
Author: Geert Uytterhoeven <geert+renesas@glider.be>
Date:   Mon Jul 17 15:25:40 2023 +0200

    drm/armada: Fix off-by-one error in armada_overlay_get_property()

    [ Upstream commit 5f0d984053f74983a287100a9519b2fabb785fb5 ]

    As ffs() returns one more than the index of the first bit set (zero
    means no bits set), the color key mode value is shifted one position too
    much.

    Fix this by using FIELD_GET() instead.

    Fixes: c96103b6c49ff9a8 ("drm/armada: move colorkey properties into overlay plane
state")
    Signed-off-by: Geert Uytterhoeven <geert+renesas@glider.be>
    Reviewed-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: Javier Martinez Canillas <javierm@redhat.com>
```

```
    Link:
https://patchwork.freedesktop.org/patch/msgid/a4d779d954a7515ddbbf31cb0f0d8184c0e7c879.16896
00265.git.geert+renesas@glider.be
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 61db62b7c3fe150b5e79bcc3b3c41c31c8628ff5
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Sun Jul 23 21:54:14 2023 +0200

    ARM: dts: BCM53573: Fix Tenda AC9 switch CPU port

    [ Upstream commit 7141209db9c335ab261a17933809a3e660ebdc12 ]

    Primary Ethernet interface is connected to the port 8 (not 5).

    Fixes: 64612828628c ("ARM: dts: BCM53573: Add Tenda AC9 switch ports")
    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Link: https://lore.kernel.org/r/20230723195416.7831-1-zajec5@gmail.com
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2fb9667c180d621b2485deff21c0c24f7992c544
Author: Jocelyn Falempe <jfalempe@redhat.com>
Date:   Thu Jul 13 15:41:31 2023 +0200

    drm/ast: report connection status on Display Port.

    [ Upstream commit f81bb0ac7872893241319ea82504956676ef02fd ]

    Aspeed always report the display port as "connected", because it
    doesn't set a .detect_ctx callback.
    Fix this by providing the proper detect callback for astdp and dp501.

    This also fixes the following regression:
    Since commit fae7d186403e ("drm/probe-helper: Default to 640x480 if no
    EDID on DP") The default resolution is now 640x480 when no monitor is
    connected. But Aspeed graphics is mostly used in servers, where no monitor
    is attached. This also affects the remote BMC resolution to 640x480, which
    is inconvenient, and breaks the anaconda installer.

    v2: Add .detect callback to the dp/dp501 connector (Jani Nikula)
    v3: Use .detect_ctx callback, and refactors (Thomas Zimmermann)
        Add a BMC virtual connector
    v4: Better indent detect_ctx() functions (Thomas Zimmermann)
    v5: Enable polling of the dp and dp501 connector status
        (Thomas Zimmermann)
    v6: Change check order in ast_astdp_is_connected (Jammy Huang)

    Fixes: fae7d186403e ("drm/probe-helper: Default to 640x480 if no EDID on DP")
    Signed-off-by: Jocelyn Falempe <jfalempe@redhat.com>
    Reviewed-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230713134316.332502-2-
jfalempe@redhat.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00214567de90f01d90735a951a16d0736dc31d85
Author: Zeyan Li <qaz6750@outlook.com>
Date:   Thu Jul 27 10:53:21 2023 +0800

    arm64: dts: qcom: sm8150: Fix the I2C7 interrupt

    [ Upstream commit f9568d22ce06192a7e14bda3a29dc216659554ff ]

    I2C6 and I2C7 use the same interrupts, which is incorrect.
    In the downstream kernel, I2C7 has interrupts of 608 instead of 607.
```

```
    Fixes: 81bee6953b58 ("arm64: dts: qcom: sm8150: add i2c nodes")
    Signed-off-by: Zeyan Li <qaz6750@outlook.com>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link:
https://lore.kernel.org/r/SY7P282MB378712225CBCEA95FE71554DB201A@SY7P282MB3787.AUSP282.PROD.
OUTLOOK.COM
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2dd8ee9de71ad8447f8459fb01dade7f6c7132da
Author: Ruan Jinjie <ruanjinjie@huawei.com>
Date:   Thu Jul 27 16:02:46 2023 +0800

    of: unittest: fix null pointer dereferencing in of_unittest_find_node_by_name()

    [ Upstream commit d6ce4f0ea19c32f10867ed93d8386924326ab474 ]

    when kmalloc() fail to allocate memory in kasprintf(), name
    or full_name will be NULL, strcmp() will cause
    null pointer dereference.

    Fixes: 0d638a07d3a1 ("of: Convert to using %pOF instead of full_name")
    Signed-off-by: Ruan Jinjie <ruanjinjie@huawei.com>
    Link: https://lore.kernel.org/r/20230727080246.519539-1-ruanjinjie@huawei.com
    Signed-off-by: Rob Herring <robh@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1030916de8d9730a88f9529fdd390e90a98f4259
Author: Yangtao Li <frank.li@vivo.com>
Date:   Mon Jul 10 11:23:49 2023 +0800

    drm/tegra: dpaux: Fix incorrect return value of platform_get_irq

    [ Upstream commit 2a1ca44b654346cadfc538c4fb32eecd8daf3140 ]

    When platform_get_irq fails, we should return dpaux->irq
    instead of -ENXIO.

    Fixes: 6b6b604215c6 ("drm/tegra: Add eDP support")
    Signed-off-by: Yangtao Li <frank.li@vivo.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230710032355.72914-13-
frank.li@vivo.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d436dcacf98e03d70c439a11b05be8f303ef2fa
Author: Ryan McCann <quic_rmccann@quicinc.com>
Date:   Fri Jul 7 18:24:40 2023 -0700

    drm/msm: Update dev core dump to not print backwards

    [ Upstream commit 903705111d863ed8ccf73465da77d232fc422ec1 ]

    Device core dump add block method adds hardware blocks to dumping queue
    with stack behavior which causes the hardware blocks to be printed in
    reverse order. Change the addition to dumping queue data structure
    from "list_add" to "list_add_tail" for FIFO queue behavior.

    Fixes: 98659487b845 ("drm/msm: add support to take dpu snapshot")
    Reviewed-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Abhinav Kumar <quic_abhinavk@quicinc.com>
    Signed-off-by: Ryan McCann <quic_rmccann@quicinc.com>
    Patchwork: https://patchwork.freedesktop.org/patch/546200/
    Link: https://lore.kernel.org/r/20230622-devcoredump_patch-v5-1-67e8b66c4723@quicinc.com
```

```
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69b0a5341457edd5baef972dc64797b89f3a423b
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Thu Jul 6 16:37:27 2023 +0800

    md/md-bitmap: hold 'reconfig_mutex' in backlog_store()

    [ Upstream commit 44abfa6a95df425c0660d56043020b67e6d93ab8 ]

    Several reasons why 'reconfig_mutex' should be held:

    1) rdev_for_each() is not safe to be called without the lock, because
       rdev can be removed concurrently.
    2) mddev_destroy_serial_pool() and mddev_create_serial_pool() should not
       be called concurrently.
    3) mddev_suspend() from mddev_destroy/create_serial_pool() should be
       protected by the lock.

    Fixes: 10c92fca636e ("md-bitmap: create and destroy wb_info_pool with the change of
backlog")
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230706083727.608914-3-yukuai1@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1958858bc85503f44f621f65095fbaadd48be919
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Thu Jul 6 16:37:26 2023 +0800

    md/md-bitmap: remove unnecessary local variable in backlog_store()

    [ Upstream commit b4d129640f194ffc4cc64c3e97f98ae944c072e8 ]

    Local variable is definied first in the beginning of backlog_store(),
    there is no need to define it again.

    Fixes: 8c13ab115b57 ("md/bitmap: don't set max_write_behind if there is no write mostly
device")
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230706083727.608914-2-yukuai1@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c3e9a852ff58997cdc674fd4d12a705bd3ca17e3
Author: Li Nan <linan122@huawei.com>
Date:   Sat Jul 1 16:05:29 2023 +0800

    md/raid10: use dereference_rdev_and_rrdev() to get devices

    [ Upstream commit 673643490b9a0eb3b25633abe604f62b8f63dba1 ]

    Commit 2ae6aaf76912 ("md/raid10: fix io loss while replacement replace
    rdev") reads replacement first to prevent io loss. However, there are same
    issue in wait_blocked_dev() and raid10_handle_discard(), too. Fix it by
    using dereference_rdev_and_rrdev() to get devices.

    Fixes: d30588b2731f ("md/raid10: improve raid10 discard request")
    Fixes: f2e7e269a752 ("md/raid10: pull the code that wait for blocked dev into one
function")
    Signed-off-by: Li Nan <linan122@huawei.com>
    Link: https://lore.kernel.org/r/20230701080529.2684932-4-linan666@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit b547a580eea8527c9a7eb9b2409ec263a8ba77f1
Author: Li Nan <linan122@huawei.com>
Date:   Sat Jul 1 16:05:28 2023 +0800

    md/raid10: factor out dereference_rdev_and_rrdev()

    [ Upstream commit b99f8fd2d91eb734f13098aa1cf337edaca454b7 ]

    Factor out a helper to get 'rdev' and 'replacement' from config->mirrors.
    Just to make code cleaner and prepare to fix the bug of io loss while
    'replacement' replace 'rdev'.

    There is no functional change.

    Signed-off-by: Li Nan <linan122@huawei.com>
    Link: https://lore.kernel.org/r/20230701080529.2684932-3-linan666@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Stable-dep-of: 673643490b9a ("md/raid10: use dereference_rdev_and_rrdev() to get
devices")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 189b8da66528d566b2f17539e23587c5d32a3f3e
Author: Yu Kuai <yukuai3@huawei.com>
Date:   Wed Jun 28 09:29:31 2023 +0800

    md: restore 'noio_flag' for the last mddev_resume()

    [ Upstream commit e24ed04389f9619e0aaef615a8948633c182a8b0 ]

    memalloc_noio_save() is called for the first mddev_suspend(), and
    repeated mddev_suspend() only increase 'suspended'. However,
    memalloc_noio_restore() is also called for the first mddev_resume(),
    which means that memory reclaim will be enabled before the last
    mddev_resume() is called, while the array is still suspended.

    Fix this problem by restore 'noio_flag' for the last mddev_resume().

    Fixes: 78f57ef9d50a ("md: use memalloc scope APIs in mddev_suspend()/mddev_resume()")
    Signed-off-by: Yu Kuai <yukuai3@huawei.com>
    Link: https://lore.kernel.org/r/20230628012931.88911-3-yukuai1@huaweicloud.com
    Signed-off-by: Song Liu <song@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4e2ce936d7c6cc84983c541e7711e9c05a948a10
Author: Herve Codina <herve.codina@bootlin.com>
Date:   Wed Jul 26 18:16:20 2023 +0200

    ASoC: fsl: fsl_qmc_audio: Fix snd_pcm_format_t values handling

    [ Upstream commit 5befe22b3eebd07b334b2917f6d14ce7ee4c8404 ]

    Running sparse on fsl_qmc_audio (make C=1) raises the following warnings:
     fsl_qmc_audio.c:387:26: warning: restricted snd_pcm_format_t degrades to integer
     fsl_qmc_audio.c:389:59: warning: incorrect type in argument 1 (different base types)
     fsl_qmc_audio.c:389:59:    expected restricted snd_pcm_format_t [usertype] format
     fsl_qmc_audio.c:389:59:    got unsigned int [assigned] i
     fsl_qmc_audio.c:564:26: warning: restricted snd_pcm_format_t degrades to integer
     fsl_qmc_audio.c:569:50: warning: incorrect type in argument 1 (different base types)
     fsl_qmc_audio.c:569:50:    expected restricted snd_pcm_format_t [usertype] format
     fsl_qmc_audio.c:569:50:    got int [assigned] i
     fsl_qmc_audio.c:573:62: warning: incorrect type in argument 1 (different base types)
     fsl_qmc_audio.c:573:62:    expected restricted snd_pcm_format_t [usertype] format
     fsl_qmc_audio.c:573:62:    got int [assigned] i
```

```
    These warnings are due to snd_pcm_format_t values handling done in the
    driver. Some macros and functions exist to handle safely these values.

    Use dedicated macros and functions to remove these warnings.

    Fixes: 075c7125b11c ("ASoC: fsl: Add support for QMC audio")
    Signed-off-by: Herve Codina <herve.codina@bootlin.com>
    Link: https://lore.kernel.org/r/20230726161620.495298-1-herve.codina@bootlin.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a7c1cbf212cf17ae5b990a809216c02fae3a50f5
Author: YingKun Meng <mengyingkun@loongson.cn>
Date:   Wed Jul 26 19:05:16 2023 +0800

    ASoC: loongson: drop of_match_ptr for OF device id

    [ Upstream commit c17bd30d0ba5ca59266771cdfc387f26271a7042 ]

    The ASoC Sound Card driver can be compile tested with !CONFIG_OF
    making 'loongson_asoc_dt_ids' unused:

    sound/soc/loongson/loongson_card.c:200:34: warning: unused variable
'loongson_asoc_dt_ids' [-Wunused-const-variable]

    As krzysztof advice, we drop of_match_ptr so the device id
    can also be used on ACPI.

    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202307242008.xqdjgk04-lkp@intel.com
    Fixes: d24028606e76 ("ASoC: loongson: Add Loongson ASoC Sound Card Support")
    Signed-off-by: YingKun Meng <mengyingkun@loongson.cn>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230726110516.703342-1-mengyingkun@loongson.cn
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ca5d032203e01f6f40f3e6f9b9ec6e31a027fcff
Author: Jonas Karlman <jonas@kwiboo.se>
Date:   Mon Jul 24 14:52:16 2023 +0000

    arm64: dts: rockchip: Enable SATA on Radxa E25

    [ Upstream commit 2bdfe84fbd57a4ed9fd65a67210442559ce078f0 ]

    The M.2 KEY B port can be used for WWAN USB2 modules or SATA drives.

    Enable sata1 node to fix use of SATA drives on the M.2 slot.

    Fixes: 2bf2f4d9f673 ("arm64: dts: rockchip: Add Radxa CM3I E25")
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Link: https://lore.kernel.org/r/20230724145213.3833099-1-jonas@kwiboo.se
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 46599192a20f3a109e5f6bb2e1c26be4b2d24ff4
Author: Jonas Karlman <jonas@kwiboo.se>
Date:   Mon Jul 24 14:52:16 2023 +0000

    arm64: dts: rockchip: Fix PCIe regulators on Radxa E25

    [ Upstream commit a87852e37f782257ebc57cc44a0d3fbf806471f6 ]

    Despite its name, the regulator vcc3v3_pcie30x1 has nothing to do with
    pcie30x1. Instead, it supply power to VBAT1-5 on the M.2 KEY B port as
```

```
    seen on page 8 of the schematic [1].

    pcie30x1 is used for the mini PCIe slot, and as seen on page 9 the
    vcc3v3_minipcie regulator is instead related to pcie30x1.

    The M.2 KEY B port can be used for WWAN USB2 modules or SATA drives.

    Use correct regulator vcc3v3_minipcie for pcie30x1.

    [1] https://dl.radxa.com/cm3p/e25/radxa-e25-v1.4-sch.pdf

    Fixes: 2bf2f4d9f673 ("arm64: dts: rockchip: Add Radxa CM3I E25")
    Signed-off-by: Jonas Karlman <jonas@kwiboo.se>
    Link: https://lore.kernel.org/r/20230724145213.3833099-1-jonas@kwiboo.se
    Signed-off-by: Heiko Stuebner <heiko@sntech.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc9d0d2aa32a247ae9f51f57c57ef8907846ab1f
Author: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
Date:   Sun Jul 23 12:19:26 2023 +0530

    drm/amdgpu: Use seq_puts() instead of seq_printf()

    [ Upstream commit fc8e55f378cf11f3abe25ec5cd67b6fc5e915a96 ]

    For a constant format without additional arguments, use seq_puts()
    instead of seq_printf(). Also, it fixes the following warning.

    WARNING: Prefer seq_puts to seq_printf

    And other style fixes:

    WARNING: Prefer 'unsigned int' to bare use of 'unsigned'
    WARNING: Block comments should align the * on each line

    Cc: Christian König <christian.koenig@amd.com>
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b7292564c153064cffc0fdad9c955531b9f834bc
Author: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
Date:   Sun Jul 23 12:29:14 2023 +0530

    drm/amdgpu: Update min() to min_t() in 'amdgpu_info_ioctl'

    [ Upstream commit a0cc8e1512ad72c9f97cdcb76d42715730adaf62 ]

    Fixes the following:

    WARNING: min() should probably be min_t(size_t, size, sizeof(ip))
    +               ret = copy_to_user(out, &ip, min((size_t)size, sizeof(ip)));

    And other style fixes:

    WARNING: Prefer 'unsigned int' to bare use of 'unsigned'
    WARNING: Missing a blank line after declarations

    Cc: Christian König <christian.koenig@amd.com>
    Cc: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Srinivasan Shanmugam <srinivasan.shanmugam@amd.com>
    Reviewed-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
```

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

commit 85ef84429eeeaf40892934777f7d4852aba3505c
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 27 18:24:25 2023 +0200

    dt-bindings: arm: msm: kpss-acc: Make the optional reg truly optional

    [ Upstream commit 7dc3ea5ea8e8df2a82a1e78bef2382fb2c982ed3 ]

    The description of reg[1] says that register is optional. Adjust
    minItems to make it truly optional.

    Fixes: 4260ddfb6496 ("dt-bindings: arm: msm: Convert and split kpss-acc driver
Documentation to yaml")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Acked-by: Rob Herring <robh@kernel.org>
    Link: https://lore.kernel.org/r/20230627-topic-more_bindings-v1-9-
6b4b6cd081e5@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0b177925a01be0486326766674c2d48e28b7cd3c
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 10:35:00 2023 +0200

    arm64: dts: qcom: sc8180x-pmics: align LPG node name with dtschema

    [ Upstream commit 4af302a7e29e70bd930e80ab8f967da48a99a31a ]

    Bindings expect the LPG node name to be "pwm":

      sc8180x-lenovo-flex-5g.dtb: pmic@5: 'lpg' does not match any of the regexes:

    Fixes: d3302290f59e ("arm64: dts: qcom: sc8180x: Add pmics")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230720083500.73554-4-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 49acafbe355e16d39e59b6b450f6edb57151db5d
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 10:34:59 2023 +0200

    arm64: dts: qcom: sc8180x-pmics: align SPMI PMIC Power-on node name with dtschema

    [ Upstream commit bf520227bd32381c587fa36271475e035daab3d7 ]

    Bindings expect the Power-on node name to be "pon":

      sc8180x-lenovo-flex-5g.dtb: pmic@0: 'power-on@800' does not match any of the regexes:

    Fixes: d3302290f59e ("arm64: dts: qcom: sc8180x: Add pmics")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230720083500.73554-3-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 47794ac7285901e7a07d14f5eb5dfb7f4a07a725
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 10:34:58 2023 +0200
```

```
    arm64: dts: qcom: sc8180x-pmics: add missing gpio-ranges

    [ Upstream commit 565951b1202e1984154abaae4567f16f8073fca3 ]

    The GPIO children of PMICs should use gpio-ranges:

      sc8180x-primus.dtb: pmic@0: gpio@c000: 'gpio-ranges' is a required property

    Fixes: d3302290f59e ("arm64: dts: qcom: sc8180x: Add pmics")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230720083500.73554-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7e2229d14234bbea8fbb5e426d5f3533b0f1b262
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 10:34:57 2023 +0200

    arm64: dts: qcom: sc8180x-pmics: add missing qcom,spmi-gpio fallbacks

    [ Upstream commit 0304fc1de3d930db83749cca6ccb3a4f89918fc4 ]

    The GPIO children of PMICs should use qcom,spmi-gpio fallback:

      sc8180x-primus.dtb: pmic@0: gpio@c000:compatible: ['qcom,pmc8180-gpio'] is too short

    Fixes: d3302290f59e ("arm64: dts: qcom: sc8180x: Add pmics")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Vinod Koul <vkoul@kernel.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230720083500.73554-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 00f3fb37f31d320b80f520811f24951ba2de2eaa
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Thu Jul 20 13:53:31 2023 +0200

    arm64: dts: qcom: msm8996-gemini: fix touchscreen VIO supply

    [ Upstream commit 21fc24ee9c5943732c9ae538766c9be93d70d936 ]

    According to bindings and Linux driver, there is no VDDA but VIO supply.

    Fixes: 4ac46b3682c5 ("arm64: dts: qcom: msm8996: xiaomi-gemini: Add support for Xiaomi
Mi 5")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230720115335.137354-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8566e038d9d5959c1603053c512dcaf179f6a598
Author: Bjorn Andersson <quic_bjorande@quicinc.com>
Date:   Mon Jun 12 15:06:32 2023 -0700

    arm64: dts: qcom: sc8180x: Fix LLCC reg property

    [ Upstream commit 74cf6675c35ec3034053a69926f4d98e52852eb0 ]

    The LLCC binding and driver was recently corrected to handle the stride
    varying between platforms. Switch to the new format to ensure accesses
    are done in the right place.
```

```
    Fixes: 8575f197b077 ("arm64: dts: qcom: Introduce the SC8180x platform")
    Signed-off-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Link: https://lore.kernel.org/r/20230612220632.1885175-1-quic_bjorande@quicinc.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bb9632703b231d4c3efca00d67975d8a164ad185
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Thu Jul 20 11:10:49 2023 +0530

    arm64: dts: qcom: sdm845: Fix the min frequency of "ice_core_clk"

    [ Upstream commit bbbef6e24bc4493602df68b052f6f48d48e3184a ]

    Minimum frequency of the "ice_core_clk" should be 75MHz as specified in the
    downstream vendor devicetree. So fix it!

    https://git.codelinaro.org/clo/la/kernel/msm-4.9/-/blob/LA.UM.7.3.r1-09300-
sdm845.0/arch/arm64/boot/dts/qcom/sdm845.dtsi

    Fixes: 433f9a57298f ("arm64: dts: sdm845: add Inline Crypto Engine registers and clock")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20230720054100.9940-5-manivannan.sadhasivam@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f9527aa02771d9cb814d856a81c51bd999544693
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Thu Jul 20 11:10:48 2023 +0530

    arm64: dts: qcom: sdm845: Add missing RPMh power domain to GCC

    [ Upstream commit 4b6ea15c0a1122422b44bf6c47a3c22fc8d46777 ]

    GCC and it's GDSCs are under the RPMh CX power domain. So let's add the
    missing RPMh power domain to the GCC node.

    Fixes: 6d4cf750d03a ("arm64: dts: sdm845: Add minimal dts/dtsi files for sdm845 SoC and
MTP")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Co-developed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Link: https://lore.kernel.org/r/20230720054100.9940-4-manivannan.sadhasivam@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 30194d0c65cdae428fd178fbcc5dfde90f211998
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Thu Jul 13 13:11:45 2023 +0200

    ARM: dts: BCM53573: Fix Ethernet info for Luxul devices

    [ Upstream commit 44ad8207806973f4e4f7d870fff36cc01f494250 ]

    Both Luxul's XAP devices (XAP-810 and XAP-1440) are access points that
    use a non-default design. They don't include switch but have a single
    Ethernet port and BCM54210E PHY connected to the Ethernet controller's
    MDIO bus.

    Support for those devices regressed due to two changes:

    1. Describing MDIO bus with switch
    After commit 9fb90ae6cae7 ("ARM: dts: BCM53573: Describe on-SoC BCM53125
```

```
      rev 4 switch") Linux stopped probing for MDIO devices.

      2. Dropping hardcoded BCM54210E delays
      In commit fea7fda7f50a ("net: phy: broadcom: Fix RGMII delays
      configuration for BCM54210E") support for other PHY modes was added but
      that requires a proper "phy-mode" value in DT.

      Both above changes are correct (they don't need to be reverted or
      anything) but they need this fix for DT data to be correct and for Linux
      to work properly.

      Fixes: 9fb90ae6cae7 ("ARM: dts: BCM53573: Describe on-SoC BCM53125 rev 4 switch")
      Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
      Link: https://lore.kernel.org/r/20230713111145.14864-1-zajec5@gmail.com
      Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit dd57f89005ca64a62a912f779c57d14c3648d1f5
  Author: Bogdan Togorean <bogdan.togorean@analog.com>
  Date:   Wed Jul 19 09:01:43 2023 +0300

      drm: adv7511: Fix low refresh rate register for ADV7533/5

      [ Upstream commit d281eeaa4de2636ff0c8e6ae387bb07b50e5fcbb ]

      For ADV7533 and ADV7535 low refresh rate is selected using
      bits [3:2] of 0x4a main register.
      So depending on ADV model write 0xfb or 0x4a register.

      Fixes: 2437e7cd88e8 ("drm/bridge: adv7533: Initial support for ADV7533")
      Reviewed-by: Robert Foss <rfoss@kernel.org>
      Reviewed-by: Nuno Sa <nuno.sa@analog.com>
      Signed-off-by: Bogdan Togorean <bogdan.togorean@analog.com>
      Signed-off-by: Alexandru Ardelean <alex@shruggie.ro>
      Reviewed-by: Frieder Schrempf <frieder.schrempf@kontron.de>
      Signed-off-by: Robert Foss <rfoss@kernel.org>
      Link: https://patchwork.freedesktop.org/patch/msgid/20230719060143.63649-1-
  alex@shruggie.ro
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 9812aa30dd6f217f5d461a805080d66c168e8deb
  Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
  Date:   Thu Jul 13 17:29:26 2023 +0200

      ARM: dts: samsung: s5pv210-smdkv210: correct ethernet reg addresses (split)

      [ Upstream commit 982655cb0e7f18934d7532c32366e574ad61dbd7 ]

      The davicom,dm9000 Ethernet Controller accepts two reg addresses.

      Fixes: b672b27d232e ("ARM: dts: Add Device tree for s5pc110/s5pv210 boards")
      Reviewed-by: Alim Akhtar <alim.akhtar@samsung.com>
      Link: https://lore.kernel.org/r/20230713152926.82884-2-krzysztof.kozlowski@linaro.org
      Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit ad1cb1fc321f692c3cfbf5e1135e7e29d9dcc662
  Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
  Date:   Thu Jul 13 17:29:25 2023 +0200

      ARM: dts: samsung: s3c6410-mini6410: correct ethernet reg addresses (split)

      [ Upstream commit cf0cb2af6a18f28b84f9f1416bff50ca60d6e98a ]

      The davicom,dm9000 Ethernet Controller accepts two reg addresses.
```

```
    Fixes: a43736deb47d ("ARM: dts: Add dts file for S3C6410-based Mini6410 board")
    Reviewed-by: Alim Akhtar <alim.akhtar@samsung.com>
    Link: https://lore.kernel.org/r/20230713152926.82884-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2c65b5407a99ad22783c375552d2aa83065d27a6
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Mon Jul 10 17:12:01 2023 +0800

    drm/bridge: anx7625: Use common macros for HDCP capabilities

    [ Upstream commit 41639b3a8b0f1f194dfe0577d99db70613f78626 ]

    The DRM DP code has macros for the DP HDCP capabilities. Use them in the
    anx7625 driver instead of raw numbers.

    Fixes: cd1637c7e480 ("drm/bridge: anx7625: add HDCP support")
    Suggested-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: Robert Foss <rfoss@kernel.org>
    Signed-off-by: Robert Foss <rfoss@kernel.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230710091203.1874317-1-
wenst@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 13acfe30ea2c480e68d92854b9580522bca07bda
Author: Chen-Yu Tsai <wenst@chromium.org>
Date:   Mon Jul 10 17:09:27 2023 +0800

    drm/bridge: anx7625: Use common macros for DP power sequencing commands

    [ Upstream commit 2ba776f903cb7157e80b5f314fb0b4faf6ea6958 ]

    The DRM DP code has macros for the DP power sequencing commands. Use
    them in the anx7625 driver instead of raw numbers.

    Fixes: 548b512e144f ("drm/bridge: anx7625: send DPCD command to downstream")
    Fixes: 27f26359de9b ("drm/bridge: anx7625: Set downstream sink into normal status")
    Signed-off-by: Chen-Yu Tsai <wenst@chromium.org>
    Reviewed-by: Nícolas F. R. A. Prado <nfraprado@collabora.com>
    Signed-off-by: Robert Foss <rfoss@kernel.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230710090929.1873646-1-
wenst@chromium.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a06070ac94f98dadb176a226dfaf3f9b66ed7614
Author: Janusz Krzysztofik <janusz.krzysztofik@linux.intel.com>
Date:   Mon Jul 10 09:36:14 2023 +0200

    x86/mm: Fix PAT bit missing from page protection modify mask

    [ Upstream commit 548cb932051fb6232ac983ed6673dae7bdf3cf4c ]

    Visible glitches have been observed when running graphics applications on
    Linux under Xen hypervisor.  Those observations have been confirmed with
    failures from kms_pwrite_crc Intel GPU test that verifies data coherency
    of DRM frame buffer objects using hardware CRC checksums calculated by
    display controllers, exposed to userspace via debugfs.  Affected
    processing paths have then been identified with new IGT test variants that
    mmap the objects using different methods and caching modes [1].

    When running as a Xen PV guest, Linux uses Xen provided PAT configuration
    which is different from its native one.  In particular, Xen specific PTE
```

```
    encoding of write-combining caching, likely used by graphics applications,
    differs from the Linux default one found among statically defined minimal
    set of supported modes.  Since Xen defines PTE encoding of the WC mode as
    _PAGE_PAT, it no longer belongs to the minimal set, depends on correct
    handling of _PAGE_PAT bit, and can be mismatched with write-back caching.

    When a user calls mmap() for a DRM buffer object, DRM device specific
    .mmap file operation, called from mmap_region(), takes care of setting PTE
    encoding bits in a vm_page_prot field of an associated virtual memory area
    structure.  Unfortunately, _PAGE_PAT bit is not preserved when the vma's
    .vm_flags are then applied to .vm_page_prot via vm_set_page_prot().  Bits
    to be preserved are determined with _PAGE_CHG_MASK symbol that doesn't
    cover _PAGE_PAT.  As a consequence, WB caching is requested instead of WC
    when running under Xen (also, WP is silently changed to WT, and UC
    downgraded to UC_MINUS).  When running on bare metal, WC is not affected,
    but WP and WT extra modes are unintentionally replaced with WC and UC,
    respectively.

    WP and WT modes, encoded with _PAGE_PAT bit set, were introduced by commit
    281d4078bec3 ("x86: Make page cache mode a real type").  Care was taken
    to extend _PAGE_CACHE_MASK symbol with that additional bit, but that
    symbol has never been used for identification of bits preserved when
    applying page protection flags.  Support for all cache modes under Xen,
    including the problematic WC mode, was then introduced by commit
    47591df50512 ("xen: Support Xen pv-domains using PAT").

    The issue needs to be fixed by including _PAGE_PAT bit into a bitmask used
    by pgprot_modify() for selecting bits to be preserved.  We can do that
    either internally to pgprot_modify() (as initially proposed), or by making
    _PAGE_PAT a part of _PAGE_CHG_MASK.  If we go for the latter then, since
    _PAGE_PAT is the same as _PAGE_PSE, we need to note that _HPAGE_CHG_MASK
    -- a huge pmds' counterpart of _PAGE_CHG_MASK, introduced by commit
    c489f1257b8c ("thp: add pmd_modify"), defined as (_PAGE_CHG_MASK |
    _PAGE_PSE) -- will no longer differ from _PAGE_CHG_MASK.  If such
    modification of _PAGE_CHG_MASK was irrelevant to its users then one might
    wonder why that new _HPAGE_CHG_MASK symbol was introduced instead of
    reusing the existing one with that otherwise irrelevant bit (_PAGE_PSE in
    that case) added.

    Add _PAGE_PAT to _PAGE_CHG_MASK and _PAGE_PAT_LARGE to _HPAGE_CHG_MASK for
    symmetry.  Split out common bits from both symbols to a common symbol for
    clarity.

    [ dhansen: tweak the solution changelog description ]

    [1] https://gitlab.freedesktop.org/drm/igt-gpu-tools/-/commit/0f0754413f14

    Fixes: 281d4078bec3 ("x86: Make page cache mode a real type")
    Signed-off-by: Janusz Krzysztofik <janusz.krzysztofik@linux.intel.com>
    Signed-off-by: Dave Hansen <dave.hansen@linux.intel.com>
    Reviewed-by: Andi Shyti <andi.shyti@linux.intel.com>
    Reviewed-by: Juergen Gross <jgross@suse.com>
    Tested-by: Marek Marczykowski-Górecki <marmarek@invisiblethingslab.com>
    Link: https://gitlab.freedesktop.org/drm/intel/-/issues/7648
    Link: https://lore.kernel.org/all/20230710073613.8006-2-
janusz.krzysztofik%40linux.intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1e4e8081edb1c154aedab9d9b1089d950656084f
Author: Chengming Zhou <zhouchengming@bytedance.com>
Date:   Mon Jul 17 12:00:56 2023 +0800

    blk-flush: fix rq->flush.seq for post-flush requests

    [ Upstream commit 28b241237470981a96fbd82077c8044466b61e5f ]
```

```
    If the policy == (REQ_FSEQ_DATA | REQ_FSEQ_POSTFLUSH), it means that the
    data sequence and post-flush sequence need to be done for this request.

    The rq->flush.seq should record what sequences have been done (or don't
    need to be done). So in this case, pre-flush doesn't need to be done,
    we should init rq->flush.seq to REQ_FSEQ_PREFLUSH not REQ_FSEQ_POSTFLUSH.

    Fixes: 615939a2ae73 ("blk-mq: defer to the normal submission path for post-flush
requests")
    Signed-off-by: Chengming Zhou <zhouchengming@bytedance.com>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230717040058.3993930-3-chengming.zhou@linux.dev
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8f054988894936413d71d81adb22653316227fa
Author: Christoph Hellwig <hch@lst.de>
Date:   Fri Jul 7 11:42:39 2023 +0200

    block: don't allow enabling a cache on devices that don't support it

    [ Upstream commit 43c9835b144c7ce29efe142d662529662a9eb376 ]

    Currently the write_cache attribute allows enabling the QUEUE_FLAG_WC
    flag on devices that never claimed the capability.

    Fix that by adding a QUEUE_FLAG_HW_WC flag that is set by
    blk_queue_write_cache and guards re-enabling the cache through sysfs.

    Note that any rescan that calls blk_queue_write_cache will still
    re-enable the write cache as in the current code.

    Fixes: 93e9d8e836cb ("block: add ability to flag write back caching on a device")
    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230707094239.107968-3-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 213007a3250c1acca1228038e5874ddc7b7c4ae8
Author: Christoph Hellwig <hch@lst.de>
Date:   Fri Jul 7 11:42:38 2023 +0200

    block: cleanup queue_wc_store

    [ Upstream commit c4e21bcd0f9d01f9c5d6c52007f5541871a5b1de ]

    Get rid of the local queue_wc_store variable and handling setting and
    clearing the QUEUE_FLAG_WC flag diretly instead the if / else if.

    Signed-off-by: Christoph Hellwig <hch@lst.de>
    Link: https://lore.kernel.org/r/20230707094239.107968-2-hch@lst.de
    Signed-off-by: Jens Axboe <axboe@kernel.dk>
    Stable-dep-of: 43c9835b144c ("block: don't allow enabling a cache on devices that don't
support it")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 50e54c81fd74a9ae0d69ba224b3a2a6e2b4261f5
Author: Lucas Stach <l.stach@pengutronix.de>
Date:   Fri Apr 14 16:38:10 2023 +0200

    drm/etnaviv: fix dumping of active MMU context

    [ Upstream commit 20faf2005ec85fa1a6acc9a74ff27de667f90576 ]
```

```
    gpu->mmu_context is the MMU context of the last job in the HW queue, which
    isn't necessarily the same as the context from the bad job. Dump the MMU
    context from the scheduler determined bad submit to make it work as intended.

    Fixes: 17e4660ae3d7 ("drm/etnaviv: implement per-process address spaces on MMUv2")
    Signed-off-by: Lucas Stach <l.stach@pengutronix.de>
    Reviewed-by: Christian Gmeiner <christian.gmeiner@gmail.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 635679ce47d5cc84721679ea53799bdf801cd32c
Author: Diogo Ivo <diogo.ivo@tecnico.ulisboa.pt>
Date:   Fri Jul 14 11:10:17 2023 +0100

    arm64: tegra: Fix HSUART for Smaug

    [ Upstream commit 590bfe51838f6345a6a3288507661dc9b7208464 ]

    After commit 71de0a054d0e ("arm64: tegra: Drop serial clock-names and
    reset-names") was applied, the HSUART failed to probe and the following
    error is seen:

     serial-tegra 70006300.serial: Couldn't get the reset
     serial-tegra: probe of 70006300.serial failed with error -2

    Commit 71de0a054d0e ("arm64: tegra: Drop serial clock-names and
    reset-names") is correct because the "reset-names" property is not
    needed for 8250 UARTs. However, the "reset-names" is required for the
    HSUART and should have been populated as part of commit a63c0cd83720c
    ("arm64: dts: tegra: smaug: Add Bluetooth node") that enabled the HSUART
    for the Pixel C. Fix this by populating the "reset-names" property for
    the HSUART on the Pixel C.

    Fixes: a63c0cd83720 ("arm64: dts: tegra: smaug: Add Bluetooth node")
    Signed-off-by: Diogo Ivo <diogo.ivo@tecnico.ulisboa.pt>
    Reviewed-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 585c0eee76d52415900a010835504f08284c7e99
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 22:00:29 2023 +0200

    arm64: dts: qcom: sc8180x: Add missing 'cache-unified' to L3

    [ Upstream commit e4322bb818bbcd36b441de9880fa4ac911a5eb51 ]

    Add the missing property to fix the dt checker warning:

    qcom/sc8180x-primus.dtb: l3-cache: 'cache-unified' is a required property

    Fixes: 8575f197b077 ("arm64: dts: qcom: Introduce the SC8180x platform")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230626-topic-bindingsfixups-v1-7-
254ae8642e69@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fd6bd00eb37b9ece59550e50ff045247a4ce3a22
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 22:00:28 2023 +0200

    arm64: dts: qcom: pmi8994: Add missing OVP interrupt

    [ Upstream commit 8db94432690371b1736e9a2566a9b3d8a73d5a97 ]
```

```
    Add the missing OVP interrupt. This fixes the schema warning:

    wled@d800: interrupt-names: ['short'] is too short

    Fixes: 37aa540cbd30 ("arm64: dts: qcom: pmi8994: Add WLED node")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230626-topic-bindingsfixups-v1-6-
254ae8642e69@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1dee37bf33407675fe93295a01c7399cb887e8b
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 22:00:27 2023 +0200

    arm64: dts: qcom: pmi8950: Add missing OVP interrupt

    [ Upstream commit 4d77b639531fd85b84a7079c3369908dfaddf8b2 ]

    Add the missing OVP interrupt. This fixes the schema warning:

    wled@d800: interrupt-names: ['short'] is too short

    Fixes: 0d97fdf380b4 ("arm64: dts: qcom: Add configuration for PMI8950 peripheral")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230626-topic-bindingsfixups-v1-5-
254ae8642e69@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 61afc714ff2203356e61ff6a5f8ba095a1f87989
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 22:00:26 2023 +0200

    arm64: dts: qcom: pm660l: Add missing short interrupt

    [ Upstream commit 9a4ac09db3c7413e334b4abd6b2f6de8930dd781 ]

    Add the missing short interrupt. This fixes the schema warning:

    wled@d800: interrupt-names: ['ovp'] is too short

    Fixes: 7b56a804e58b ("arm64: dts: qcom: pm660l: Add WLED support")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230626-topic-bindingsfixups-v1-4-
254ae8642e69@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit fb04684797aa922af0bab5ac6ffa4cd51c6a6db0
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Mon Jun 26 22:00:25 2023 +0200

    arm64: dts: qcom: pm6150l: Add missing short interrupt

    [ Upstream commit 7e1f024ef0d1da456f61d00f01dc3287ede915b3 ]

    Add the missing short interrupt. This fixes the schema warning:

    wled@d800: interrupt-names: ['ovp'] is too short

    Fixes: fe508ced49dd ("arm64: dts: qcom: pm6150l: Add wled node")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Luca Weiss <luca@z3ntu.xyz>
```

```
    Link: https://lore.kernel.org/r/20230626-topic-bindingsfixups-v1-3-
254ae8642e69@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c25152192a3a33fc05124bb0335b74fd1bf8b0a4
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Tue Jul 11 08:30:11 2023 +0200

    arm64: dts: qcom: sm8250-sony-xperia: correct GPIO keys wakeup again

    [ Upstream commit b8fbeea0253211d97c579eae787274633d3eaf0d ]

    gpio-keys,wakeup is a deprecated property:

      m8250-sony-xperia-edo-pdx206.dtb: gpio-keys: key-camera-focus: Unevaluated properties
are not allowed ('gpio-key,wakeup' was unexpected)

    Fixes: a422c6a91a66 ("arm64: dts: qcom: sm8250-edo: Rectify gpio-keys")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230711063011.16222-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6d0d93db168aa518b117c933e9e0649286eefb5e
Author: Jon Hunter <jonathanh@nvidia.com>
Date:   Mon Jul 3 12:36:17 2023 +0100

    arm64: tegra: Fix HSUART for Jetson AGX Orin

    [ Upstream commit 861dbb2b15b1049113887fb95e856f7123eea0cc ]

    After commit 71de0a054d0e ("arm64: tegra: Drop serial clock-names and
    reset-names") was applied, the HSUART failed to probe and the following
    error is seen:

     serial-tegra 3100000.serial: Couldn't get the reset
     serial-tegra: probe of 3100000.serial failed with error -2

    Commit 71de0a054d0e ("arm64: tegra: Drop serial clock-names and
    reset-names") is correct because the "reset-names" property is not
    needed for 8250 UARTs. However, the "reset-names" is required for the
    HSUART and should have been populated as part of commit ff578db7b693
    ("arm64: tegra: Enable UART instance on 40-pin header") that
    enabled the HSUART for Jetson AGX Orin. Fix this by populating the
    "reset-names" property for the HSUART on Jetson AGX Orin.

    Fixes: ff578db7b693 ("arm64: tegra: Enable UART instance on 40-pin header")
    Signed-off-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eaf2bb67bdb6449d89064d51d6dd9fece5b395ba
Author: Jon Hunter <jonathanh@nvidia.com>
Date:   Mon Jul 3 12:35:37 2023 +0100

    arm64: tegra: Add missing alias for NVIDIA IGX Orin

    [ Upstream commit d97966df30ed8c7df0350b8ff6662e38ee88c39f ]

    The following error is seen on boot for the NVIDIA IGX Orin platform ...

     serial-tegra 3100000.serial: failed to get alias id, errno -19
```

```
    Fix this by populating the necessary alias for the serial device.

    Fixes: c95711d7dbc4 ("arm64: tegra: Add support for IGX Orin")
    Signed-off-by: Jon Hunter <jonathanh@nvidia.com>
    Signed-off-by: Thierry Reding <treding@nvidia.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 90d1333b456b4c9dafe40bdd4054b99da600694f
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Fri Jul 7 13:40:04 2023 +0200

    ARM: dts: BCM53573: Use updated "spi-gpio" binding properties

    [ Upstream commit 2c0fd6b3d0778ceab40205315ccef74568490f17 ]

    Switch away from deprecated properties.

    This fixes:
    arch/arm/boot/dts/broadcom/bcm947189acdbmr.dtb: spi: gpio-sck: False schema does not
allow [[3, 21, 0]]
            From schema: Documentation/devicetree/bindings/spi/spi-gpio.yaml
    arch/arm/boot/dts/broadcom/bcm947189acdbmr.dtb: spi: gpio-miso: False schema does not
allow [[3, 22, 0]]
            From schema: Documentation/devicetree/bindings/spi/spi-gpio.yaml
    arch/arm/boot/dts/broadcom/bcm947189acdbmr.dtb: spi: gpio-mosi: False schema does not
allow [[3, 23, 0]]
            From schema: Documentation/devicetree/bindings/spi/spi-gpio.yaml
    arch/arm/boot/dts/broadcom/bcm947189acdbmr.dtb: spi: 'sck-gpios' is a required property
            From schema: Documentation/devicetree/bindings/spi/spi-gpio.yaml
    arch/arm/boot/dts/broadcom/bcm947189acdbmr.dtb: spi: Unevaluated properties are not
allowed ('gpio-miso', 'gpio-mosi', 'gpio-sck' were unexpected)
            From schema: Documentation/devicetree/bindings/spi/spi-gpio.yaml

    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Link: https://lore.kernel.org/r/20230707114004.2740-4-zajec5@gmail.com
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d3ee03216ec0e24d903ec154b64133959cba27b2
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Fri Jul 7 13:40:03 2023 +0200

    ARM: dts: BCM53573: Add cells sizes to PCIe node

    [ Upstream commit 3392ef368d9b04622fe758b1079b512664b6110a ]

    This fixes:
    arch/arm/boot/dts/broadcom/bcm47189-luxul-xap-1440.dtb: pcie@2000: '#address-cells' is a
required property
            From schema: /lib/python3.10/site-packages/dtschema/schemas/pci/pci-bus.yaml
    arch/arm/boot/dts/broadcom/bcm47189-luxul-xap-1440.dtb: pcie@2000: '#size-cells' is a
required property
            From schema: /lib/python3.10/site-packages/dtschema/schemas/pci/pci-bus.yaml

    Two properties that need to be added later are "device_type" and
    "ranges". Adding "device_type" on its own causes a new warning and the
    value of "ranges" needs to be determined yet.

    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Link: https://lore.kernel.org/r/20230707114004.2740-3-zajec5@gmail.com
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 639af01954eb0d9647ea58047831e3e610bd67e3
Author: Rafał Miłecki <rafal@milecki.pl>
```

```
Date:   Fri Jul 7 13:40:02 2023 +0200

    ARM: dts: BCM53573: Drop nonexistent #usb-cells

    [ Upstream commit 05d2c3d552b8c92fc397377d9d1112fc58e2cd59 ]

    Such property simply doesn't exist (is not documented or used anywhere).

    This fixes:
    arch/arm/boot/dts/broadcom/bcm47189-luxul-xap-1440.dtb: usb@d000: Unevaluated properties
are not allowed ('#usb-cells' was unexpected)
            From schema: Documentation/devicetree/bindings/usb/generic-ohci.yaml

    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Link: https://lore.kernel.org/r/20230707114004.2740-2-zajec5@gmail.com
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6657bd418423516c72ca4cce03aa2a6627656371
Author: Rafał Miłecki <rafal@milecki.pl>
Date:   Fri Jul 7 13:40:01 2023 +0200

    ARM: dts: BCM53573: Drop nonexistent "default-off" LED trigger

    [ Upstream commit be7e1e5b0f67c58ec4be0a54db23b6a4fa6e2116 ]

    There is no such trigger documented or implemented in Linux. It was a
    copy & paste mistake.

    This fixes:
    arch/arm/boot/dts/broadcom/bcm47189-luxul-xap-1440.dtb: leds: led-wlan:linux,default-
trigger: 'oneOf' conditional failed, one must be fixed:
            'default-off' is not one of ['backlight', 'default-on', 'heartbeat', 'disk-
activity', 'disk-read', 'disk-write', 'timer', 'pattern', 'audio-micmute', 'audio-mute',
'bluetooth-power', 'flash', 'kbd-capslock', 'mtd', 'nand-disk', 'none', 'torch', 'usb-
gadget', 'usb-host', 'usbport']
            'default-off' does not match '^cpu[0-9]*$'
            'default-off' does not match '^hci[0-9]+-power$'
            'default-off' does not match '^mmc[0-9]+$'
            'default-off' does not match '^phy[0-9]+tx$'
            From schema: Documentation/devicetree/bindings/leds/leds-gpio.yaml

    Signed-off-by: Rafał Miłecki <rafal@milecki.pl>
    Link: https://lore.kernel.org/r/20230707114004.2740-1-zajec5@gmail.com
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bc0103dc54f46041f4080880f39a63b6f6917493
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Fri Jul 7 13:11:51 2023 +0200

    drm/amdgpu: avoid integer overflow warning in amdgpu_device_resize_fb_bar()

    [ Upstream commit 822130b5e8834ab30ad410cf19a582e5014b9a85 ]

    On 32-bit architectures comparing a resource against a value larger than
    U32_MAX can cause a warning:

    drivers/gpu/drm/amd/amdgpu/amdgpu_device.c:1344:18: error: result of comparison of
constant 4294967296 with expression of type 'resource_size_t' (aka 'unsigned int') is always
false [-Werror,-Wtautological-constant-out-of-range-compare]
                        res->start > 0x100000000ull)
                        ~~~~~~~~~~ ^ ~~~~~~~~~~~~~~

    As gcc does not warn about this in dead code, add an IS_ENABLED() check at
```

```
    the start of the function. This will always return success but not actually resize
    the BAR on 32-bit architectures without high memory, which is exactly what
    we want here, as the driver can fall back to bank switching the VRAM
    access.

    Fixes: 31b8adab3247 ("drm/amdgpu: require a root bus window above 4GB for BAR resize")
    Reviewed-by: Christian König <christian.koenig@amd.com>
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9d7db065a205ac902dbc0f18693b64ed34f8a719
Author: Nishanth Menon <nm@ti.com>
Date:   Tue Jun 20 08:03:29 2023 -0500

    firmware: ti_sci: Use system_state to determine polling

    [ Upstream commit 9225bcdedf16297a346082e7d23b0e8434aa98ed ]

    Commit b9e8a7d950ff ("firmware: ti_sci: Switch transport to polled
    mode during system suspend") aims to resolve issues with tisci
    operations during system suspend operation. However, the system may
    enter a no_irq stage in various other usage modes, including power-off
    and restart. To determine if polling mode is appropriate, use the
    system_state instead.

    While at this, drop the unused is_suspending state variable and
    related helpers.

    Fixes: b9e8a7d950ff ("firmware: ti_sci: Switch transport to polled mode during system
suspend")
    Reported-by: Francesco Dolcini <francesco@dolcini.it>
    Reported-by: Wadim Egorov <w.egorov@phytec.de>
    Tested-by: Francesco Dolcini <francesco.dolcini@toradex.com> # Toradex Verdin AM62
    Link: https://lore.kernel.org/r/20230620130329.4120443-1-nm@ti.com
    Closes: https://lore.kernel.org/all/ZGeHMjlnob2GFyHF@francesco-nb.int.toradex.com/
    Signed-off-by: Nishanth Menon <nm@ti.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ac81bce70c9ece040d3036682c9d159c464159f0
Author: Marek Vasut <marex@denx.de>
Date:   Tue Jul 11 15:11:21 2023 +0200

    ARM: dts: stm32: Add missing detach mailbox for DHCOR SoM

    [ Upstream commit 2f38de940f072db369edd3e6e8d82bb8f42c5c9b ]

    Add missing "detach" mailbox to this board to permit the CPU to inform
    the remote processor on a detach. This signal allows the remote processor
    firmware to stop IPC communication and to reinitialize the resources for
    a re-attach.

    Without this mailbox, detach is not possible and kernel log contains the
    following warning to, so make sure all the STM32MP15xx platform DTs are
    in sync regarding the mailboxes to fix the detach issue and the warning:
    "
    stm32-rproc 10000000.m4: mbox_request_channel_byname() could not locate channel named
"detach"
    "

    Fixes: 6257dfc1c412 ("ARM: dts: stm32: Add coprocessor detach mbox on stm32mp15x-dkx
boards")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit c8a43bef2fef62bd196bb65810e3b7f397227923
Author: Marek Vasut <marex@denx.de>
Date:   Tue Jul 11 15:09:07 2023 +0200

    ARM: dts: stm32: Add missing detach mailbox for DHCOM SoM

    [ Upstream commit deb7edbc27a6ec4d8f5edfd8519b7ed13cbd2a52 ]

    Add missing "detach" mailbox to this board to permit the CPU to inform
    the remote processor on a detach. This signal allows the remote processor
    firmware to stop IPC communication and to reinitialize the resources for
    a re-attach.

    Without this mailbox, detach is not possible and kernel log contains the
    following warning to, so make sure all the STM32MP15xx platform DTs are
    in sync regarding the mailboxes to fix the detach issue and the warning:
    "
    stm32-rproc 10000000.m4: mbox_request_channel_byname() could not locate channel named
"detach"
    "

    Fixes: 6257dfc1c412 ("ARM: dts: stm32: Add coprocessor detach mbox on stm32mp15x-dkx
boards")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d6097b1c34db01ec6f9290b2cb64e8b236874fe
Author: Marek Vasut <marex@denx.de>
Date:   Thu May 18 03:12:43 2023 +0200

    ARM: dts: stm32: Add missing detach mailbox for Odyssey SoM

    [ Upstream commit 966f04a89d77548e673de2c400abe0b2cf5c15db ]

    Add missing "detach" mailbox to this board to permit the CPU to inform
    the remote processor on a detach. This signal allows the remote processor
    firmware to stop IPC communication and to reinitialize the resources for
    a re-attach.

    Without this mailbox, detach is not possible and kernel log contains the
    following warning to, so make sure all the STM32MP15xx platform DTs are
    in sync regarding the mailboxes to fix the detach issue and the warning:
    "
    stm32-rproc 10000000.m4: mbox_request_channel_byname() could not locate channel named
"detach"
    "

    Fixes: 6257dfc1c412 ("ARM: dts: stm32: Add coprocessor detach mbox on stm32mp15x-dkx
boards")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3c0f65440df9b8dffac1285eb7e8900391c02294
Author: Marek Vasut <marex@denx.de>
Date:   Thu May 18 03:12:42 2023 +0200

    ARM: dts: stm32: Add missing detach mailbox for emtrion emSBC-Argon

    [ Upstream commit 0ee0ef38aa9f75f21b51f729dd42b2e932515188 ]

    Add missing "detach" mailbox to this board to permit the CPU to inform
    the remote processor on a detach. This signal allows the remote processor
```

```
    firmware to stop IPC communication and to reinitialize the resources for
    a re-attach.

    Without this mailbox, detach is not possible and kernel log contains the
    following warning to, so make sure all the STM32MP15xx platform DTs are
    in sync regarding the mailboxes to fix the detach issue and the warning:
    "
    stm32-rproc 10000000.m4: mbox_request_channel_byname() could not locate channel named
 "detach"
    "

    Fixes: 6257dfc1c412 ("ARM: dts: stm32: Add coprocessor detach mbox on stm32mp15x-dkx
 boards")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Signed-off-by: Alexandre Torgue <alexandre.torgue@foss.st.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6fe2cb9cc1c0d7825622f1feb0ec17e2b72aa7fc
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jul 4 14:23:18 2023 +0200

    arm64: dts: qcom: sm8250: Mark SMMUs as DMA coherent

    [ Upstream commit 4cb19bd7c6329c4702f92c6dd4e7c02eb903ca13 ]

    The SMMUs on SM8250 are cache-coherent. Mark them as such.

    Fixes: a89441fcd09d ("arm64: dts: qcom: sm8250: add apps_smmu node")
    Fixes: 04a3605b184e ("arm64: dts: qcom: add sm8250 GPU nodes")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230704-topic-8250_pcie_dmac-v1-2-
 799603a980b0@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4fb17b6fac5b72d99ca0e6ec52a091211e1e3e90
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Sun Jul 2 20:50:50 2023 +0200

    arm64: dts: qcom: minor whitespace cleanup around '='

    [ Upstream commit 934a3b4d5a2d4c265ca22d3cf471a72ec8d9ee65 ]

    The DTS code coding style expects exactly one space before and after '='
    sign.

    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Link: https://lore.kernel.org/r/20230702185051.43867-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Stable-dep-of: 4cb19bd7c632 ("arm64: dts: qcom: sm8250: Mark SMMUs as DMA coherent")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 49b69b9ea6c4d74d75845d7cf448f96e2f567786
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jul 4 14:23:17 2023 +0200

    arm64: dts: qcom: sm8250: Mark PCIe hosts as DMA coherent

    [ Upstream commit 339d38a436f30d0f874815eafc7de2257346bf26 ]

    The PCIe hosts on SM8250 are cache-coherent. Mark them as such.

    Fixes: e53bdfc00977 ("arm64: dts: qcom: sm8250: Add PCIe support")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230704-topic-8250_pcie_dmac-v1-1-
```

```
  799603a980b0@linaro.org
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 590b92023316f28aaddba04ba2f9fe702dbf9229
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
  Date:   Fri Jul 7 15:30:25 2023 +0300

      arm64: dts: qcom: sm8450-hdk: remove pmr735b PMIC inclusion

      [ Upstream commit 701b59db773730a914f1778cf2dd05e3a05c2c69 ]

      The 8450-HDK doesn't use PMR735B PMIC. Drop its inclusion to remove the
      warning during the HDK bootup.

      Fixes: 30464456a1ea ("arm64: dts: qcom: sm8450-hdk: add pmic files")
      Reviewed-by: Neil Armstrong <neil.armstrong@linaro.org>
      Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
      Link: https://lore.kernel.org/r/20230707123027.1510723-7-dmitry.baryshkov@linaro.org
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit be2456fe5d11bfdb049942f5be81cde8c6c723fd
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
  Date:   Fri Jul 7 15:30:24 2023 +0300

      arm64: dts: qcom: pmk8350: fix ADC-TM compatible string

      [ Upstream commit 435a73d7377ceb29c1a22d2711dd85c831b40c45 ]

      The commit b2de43136058 ("arm64: dts: qcom: pmk8350: Add peripherals for
      pmk8350") for the ADC TM (thermal monitoring device) have used the
      compatible string from the vendor kernel ("qcom,adc-tm7"). Use the
      proper compatible string that is defined in the upstream kernel
      ("qcom,spmi-adc-tm5-gen2").

      Fixes: b2de43136058 ("arm64: dts: qcom: pmk8350: Add peripherals for pmk8350")
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
      Link: https://lore.kernel.org/r/20230707123027.1510723-6-dmitry.baryshkov@linaro.org
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit bbca97fa20f8529fc5e0a97383762a2cdf612009
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
  Date:   Fri Jul 7 15:30:23 2023 +0300

      arm64: dts: qcom: pmr735b: fix thermal zone name

      [ Upstream commit 99f8cf491d546cd668236f573c7d846d3e94f2d6 ]

      The name of the thermal zone in pmr735b.dtsi (pmr735a-thermal) conflicts
      with the thermal zone in pmr735a.dtsi. Rename the thermal zone according
      to the chip name.

      Fixes: 6f3426b3dea4 ("arm64: dts: qcom: pmr735b: add temp sensor and thermal zone
  config")
      Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
      Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
      Link: https://lore.kernel.org/r/20230707123027.1510723-5-dmitry.baryshkov@linaro.org
      Signed-off-by: Bjorn Andersson <andersson@kernel.org>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit aea6ae83f6a61fac2487a7ed8d81b5a2a1a083fd
  Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
```

```
Date:   Fri Jul 7 15:30:22 2023 +0300

    arm64: dts: qcom: pm8350b: fix thermal zone name

    [ Upstream commit aad41d9e6c44dfe299cddab97528a5333f17bdfe ]

    The name of the thermal zone in pm8350b.dtsi (pm8350c-thermal) conflicts
    with the thermal zone in pm8350c.dtsi. Rename the thermal zone according
    to the chip name.

    Fixes: 5c1399299d9d ("arm64: dts: qcom: pm8350b: add temp sensor and thermal zone
config")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Link: https://lore.kernel.org/r/20230707123027.1510723-4-dmitry.baryshkov@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6cf49301cc1668f7eb76cb724298bda35bda0ce5
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Fri Jul 7 15:30:21 2023 +0300

    arm64: dts: qcom: pm8350: fix thermal zone name

    [ Upstream commit 64f19c06f704846db5e4885ca63c689d9bef5723 ]

    The name of the thermal zone in pm8350.dtsi (pm8350c-thermal) conflicts
    with the thermal zone in pm8350c.dtsi. Rename the thermal zone according
    to the chip name.

    Fixes: 7a79b95f4288 ("arm64: dts: qcom: pm8350: add temp sensor and thermal zone
config")
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Link: https://lore.kernel.org/r/20230707123027.1510723-3-dmitry.baryshkov@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2dea1036a6a549f73afc27e216b225bd72ab44f7
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Jul 6 18:35:37 2023 +0200

    arm64: dts: qcom: sm8350: Use proper CPU compatibles

    [ Upstream commit 4390730cc12af25f7c997f477795f5f4200149c0 ]

    The Kryo names (once again) turned out to be fake. The CPUs report:

    0x412fd050 (CA55 r2p0) (0 - 3)
    0x411fd410 (CA78 r1p1) (4 - 6)
    0x411fd440 (CX1  r1p1) (7)

    Use the compatibles that reflect that.

    Fixes: b7e8f433a673 ("arm64: dts: qcom: Add basic devicetree support for SM8350 SoC")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230706-topic-sm8350-cpu-compat-v1-1-
f8d6a1869781@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b65e97bd416665e1fa2ffdc715cd378238a4675f
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Jul 5 17:00:05 2023 +0200
```

```
    arm64: dts: qcom: sc8180x: Fix cluster PSCI suspend param

    [ Upstream commit 9c31a3f5abc9eeb6509d06041b1e5f12deb39c4d ]

    The value was copypasted from 8150, but 8180 expects a different one.
    Confirmed with both downstream device tree and Windows DSDT, not tested
    on hardware (sorry, I don't have any).

    Fix it.

    Fixes: 8575f197b077 ("arm64: dts: qcom: Introduce the SC8180x platform")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Bjorn Andersson <andersson@kernel.org>
    Link: https://lore.kernel.org/r/20230705-topic-8180_sleep-v1-1-c5dce117364e@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bfa787b1ce7267d5a10153a6a098531f55bf805a
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Jul 5 15:36:23 2023 +0200

    arm64: dts: qcom: sm8350: Add missing LMH interrupts to cpufreq

    [ Upstream commit 951151c2bb548e0f6b2c40ab4c48675f5342c914 ]

    Add the missing interrupts that communicate the hardware-managed
    throttling to Linux.

    Fixes: ccbb3abb23a5 ("arm64: dts: qcom: sm8350: Add cpufreq node")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230705-topic-sm8350_fixes-v1-3-0f69f70ccb6a@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e92b3b8815a75bec0f4e1fa8b80631a998ea27f1
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Wed Jul 5 15:36:22 2023 +0200

    arm64: dts: qcom: sm8350: Fix CPU idle state residency times

    [ Upstream commit 91ce3693e2fb685f31d39605a5ad1fbd940804da ]

    The present values look to have been copypasted from 8150 or 8180.
    Fix that.

    Fixes: 07ddb302811e ("arm64: dts: qcom: sm8350: Add CPU topology and idle-states")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230705-topic-sm8350_fixes-v1-2-0f69f70ccb6a@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5107d93e83d560aa987ad9661d9f4c589c457889
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 27 19:27:50 2023 +0200

    arm64: dts: qcom: sdm845-tama: Set serial indices and stdout-path

    [ Upstream commit 9acc60c3e2d449243e4c2126e3b56f1c4f7fd3bc ]

    UART6 is used for debug (routed via uSD pins) and UART9 is connected
    to the bluetooth chip.

    Set indexed aliases to make the GENI UART driver happy and route serial
    traffic through the debug uart by default.
```

```
    Fixes: 30a7f99befc6 ("arm64: dts: qcom: Add support for SONY Xperia XZ2 / XZ2C / XZ3
(Tama platform)")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Link: https://lore.kernel.org/r/20230627-topic-tama_uart-v1-1-0fa790248db8@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e1a77e5b87e8b1d922f88a80b5687e10abb05797
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 27 18:24:27 2023 +0200

    arm64: dts: qcom: msm8996: Add missing interrupt to the USB2 controller

    [ Upstream commit 36541089c4733355ed844c67eebd0c3936953454 ]

    The interrupt line was previously not described. Take care of that.

    Fixes: 1e39255ed29d ("arm64: dts: msm8996: Add device node for qcom,dwc3")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230627-topic-more_bindings-v1-11-
6b4b6cd081e5@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 092cfda17e46391eb6137c8ae0d8df87f6512b90
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 27 18:24:19 2023 +0200

    arm64: dts: qcom: msm8939: Add missing 'cache-unified' to L2

    [ Upstream commit 68a59251f1c590ad567ff7fd799f6634fbab6e16 ]

    Add the missing property to fix the dt checker warning:

    qcom/apq8039-t2.dtb: l2-cache: 'cache-unified' is a required property

    Fixes: 61550c6c156c ("arm64: dts: qcom: Add msm8939 SoC")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Bryan O'Donoghue <bryan.odonoghue@linaro.org>
    Link: https://lore.kernel.org/r/20230627-topic-more_bindings-v1-3-
6b4b6cd081e5@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c7ea3c4a268697a44ff12e01089522b29dc6372d
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 27 18:24:18 2023 +0200

    arm64: dts: qcom: msm8939: Drop "qcom,idle-state-spc" compatible

    [ Upstream commit 982f810fc196002808b6d4230ba8f431c993d264 ]

    As of today, the only cool and legal way to get ARM64 SMP going is
    via PSCI (or spin tables). Sadly, not all chip and device vendors were
    considerate of this in the early days of arm64. Qualcomm, for example
    reused their tried-and-true spin-up method from MSM8974 and their Krait/
    arm32 Cortex designs.

    MSM8916 supports SMP with its arm32 dt overlay, as probably could 8939.
    But the arm64 DT should not define non-PSCI SMP or CPUidle stuff.

    Drop the qcom,idle-state-spc compatible (associated with Qualcomm-specific
    CPUIdle) to make the dt checker happy:
```

```
    apq8039-t2.dtb: idle-states: cpu-sleep-0:compatible:
    ['qcom,idle-state-spc', 'arm,idle-state'] is too long

    Fixes: 61550c6c156c ("arm64: dts: qcom: Add msm8939 SoC")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Benjamin Li <benl@squareup.com>
    Link: https://lore.kernel.org/r/20230627-topic-more_bindings-v1-2-
6b4b6cd081e5@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8633571b0b62e3b685ff7c02375723bf3eb2e982
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Thu Jun 22 17:56:16 2023 +0200

    arm64: dts: qcom: sc8280xp: Add missing SCM interconnect

    [ Upstream commit 0a69ccf20b0837db857abfc94d7e3bacf1cb771b ]

    The SCM interconnect path was missing. Add it.

    Fixes: 152d1faf1e2f ("arm64: dts: qcom: add SC8280XP platform")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Reviewed-by: Johan Hovold <johan+linaro@kernel.org>
    Tested-by: Johan Hovold <johan+linaro@kernel.org>
    Link: https://lore.kernel.org/r/20230622-topic-8280scmicc-v1-2-6ef318919ea5@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 622753c7fb62699b816cf925892b2d7e14ad3c5e
Author: Bjorn Andersson <quic_bjorande@quicinc.com>
Date:   Tue Jun 20 13:39:14 2023 -0700

    arm64: dts: qcom: sc8280xp-crd: Correct vreg_misc_3p3 GPIO

    [ Upstream commit 9566b5271f68bdf6e69b7c511850e3fb75cd18be ]

    The vreg_misc_3p3 regulator is controlled by PMC8280_1 GPIO 2, not 1, on
    the CRD.

    Fixes: ccd3517faf18 ("arm64: dts: qcom: sc8280xp: Add reference device")
    Signed-off-by: Bjorn Andersson <quic_bjorande@quicinc.com>
    Reviewed-by: Johan Hovold <johan+linaro@kernel.org>
    Tested-by: Johan Hovold <johan+linaro@kernel.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230620203915.141337-1-quic_bjorande@quicinc.com
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 369002cde91864ee81cb390d1d9fea31ba24912c
Author: Konrad Dybcio <konrad.dybcio@linaro.org>
Date:   Tue Jun 20 13:05:37 2023 +0200

    arm64: dts: qcom: sm8250-edo: Rectify gpio-keys

    [ Upstream commit a422c6a91a667b309ca1a6c08b30dbfcf7d4e866 ]

    Set up the corresponding GPIOs properly and add the leftover hardware
    buttons to mark this piece of the puzzle complete.

    Fixes: 46e14907c716 ("arm64: dts: qcom: sm8250-edo: Add hardware keys")
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230614-topic-edo_pinsgpiopmic-v2-4-
```

```
 6f90bba54c53@linaro.org
     Signed-off-by: Bjorn Andersson <andersson@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 763b3cb2fface66109e0f1679989025fd1648fd3
 Author: Konrad Dybcio <konrad.dybcio@linaro.org>
 Date:   Tue Jun 20 13:05:35 2023 +0200

     arm64: dts: qcom: sm8250-edo: Add GPIO line names for PMIC GPIOs

     [ Upstream commit 6b8a63350752c6a5e4b54f2de6174084652cd3cd ]

     Sony ever so graciously provides GPIO line names in their downstream
     kernel (though sometimes they are not 100% accurate and you can judge
     that by simply looking at them and with what drivers they are used).

     Add these to the PDX203&206 DTSIs to better document the hardware.

     Diff between 203 and 206:
     pm8009_gpios
     <                           "CAM_PWR_LD_EN",
     >                           "NC",

     pm8150_gpios
     <                           "NC",
     >                           "G_ASSIST_N",
     <                           "WLC_EN_N", /* GPIO_10 */
     >                           "NC", /* GPIO_10 */
     Which is due to 5 II having an additional Google Assistant hardware
     button and 1 II having a wireless charger & different camera wiring
     to accommodate the additional 3D iToF sensor.

     Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
     Link: https://lore.kernel.org/r/20230614-topic-edo_pinsgpiopmic-v2-2-
 6f90bba54c53@linaro.org
     Signed-off-by: Bjorn Andersson <andersson@kernel.org>
     Stable-dep-of: a422c6a91a66 ("arm64: dts: qcom: sm8250-edo: Rectify gpio-keys")
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 4e286c8a48805a0c4e2790f52cc73d5837d56eff
 Author: Konrad Dybcio <konrad.dybcio@linaro.org>
 Date:   Tue Jun 20 13:05:34 2023 +0200

     arm64: dts: qcom: sm8250-edo: Add gpio line names for TLMM

     [ Upstream commit 40b398beabdfe0e9088b13976e56b1dc706fe851 ]

     Sony ever so graciously provides GPIO line names in their downstream
     kernel (though sometimes they are not 100% accurate and you can judge
     that by simply looking at them and with what drivers they are used).

     Add these to the PDX203&206 DTSIs to better document the hardware.

     Diff between 203 and 206:
     <                           "CAM_PWR_A_CS",
     >                           "FRONTC_PWR_EN",
     <                           "CAM4_MCLK",
     <                           "TOF_RST_N",
     >                           "NC",
     >                           "NC",
     <                           "WLC_I2C_SDA",
     <                           "WLC_I2C_SCL", /* GPIO_120 */
     >                           "NC",
     >                           "NC",
     <                           "WLC_INT_N",
```

```
>                           "NC",

    Which makes sense, as 203 has a 3D iToF, slightly different camera
    power wiring and WLC (WireLess Charging).

    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230614-topic-edo_pinsgpiopmic-v2-1-
6f90bba54c53@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Stable-dep-of: a422c6a91a66 ("arm64: dts: qcom: sm8250-edo: Rectify gpio-keys")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 887610c1bcf587771b7de701f108f42c1361d52d
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Sat Jun 17 19:15:28 2023 +0200

    arm64: dts: qcom: msm8916-l8150: correct light sensor VDDIO supply

    [ Upstream commit 6a541eaa6e8e5283efb993ae7a947bede8d01fa5 ]

    liteon,ltr559 light sensor takes VDDIO, not VIO, supply:

      msm8916-longcheer-l8150.dtb: light-sensor@23: 'vio-supply' does not match any of the
regexes: 'pinctrl-[0-9]+'

    Fixes: 3016af34ef8d ("arm64: dts: qcom: msm8916-longcheer-l8150: Add light and proximity
sensor")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Nikita Travkin <nikita@trvn.ru>
    Link: https://lore.kernel.org/r/20230617171541.286957-2-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3e85100673c84138c0ce394c13156afcbda70acd
Author: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
Date:   Fri Jun 16 19:49:55 2023 +0200

    arm64: dts: qcom: sm8450: correct crypto unit address

    [ Upstream commit b02966f8689795406ac210189924a8cb02a71bbe ]

    Crypto node unit address should match reg.

    Fixes: b92b0d2f7582 ("arm64: dts: qcom: sm8450: add crypto nodes")
    Signed-off-by: Krzysztof Kozlowski <krzysztof.kozlowski@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230616174955.1783652-1-krzysztof.kozlowski@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6cdc0654559a2cab131c35cbb7f0460d259558a0
Author: Vincent Guittot <vincent.guittot@linaro.org>
Date:   Thu Jun 15 17:48:52 2023 +0200

    arm64: dts: qcom: sm8250: correct dynamic power coefficients

    [ Upstream commit 775a5283c25d160b2a1359018c447bc518096547 ]

    sm8250 faces the same problem with its Energy Model as sdm845. The energy
    cost of LITTLE cores is reported to be higher than medium or big cores

    EM computes the energy with formula:

    energy = OPP's cost / maximum cpu capacity * utilization
```

```
    On v6.4-rc6 we have:
    max capacity of CPU0 = 284
    capacity of CPU0's OPP(1612800 Hz) = 253
    cost of CPU0's OPP(1612800 Hz) = 191704

    max capacity of CPU4 = 871
    capacity of CPU4's OPP(710400 Hz) = 255
    cost of CPU4's OPP(710400 Hz) = 343217

    Both OPPs have almost the same compute capacity but the estimated energy
    per unit of utilization will be estimated to:

    energy CPU0 = 191704 / 284 * 1 = 675
    energy CPU4 = 343217 / 871 * 1 = 394

    EM estimates that little CPU0 will consume 71% more than medium CPU4 for
    the same compute capacity. According to [1], little consumes 25% less than
    medium core for Coremark benchmark at those OPPs for the same duration.

    Set the dynamic-power-coefficient of CPU0-3 to 105 to fix the energy model
    for little CPUs.

    [1] https://github.com/kdrag0n/freqbench/tree/master/results/sm8250/k30s

    Fixes: 6aabed5526ee ("arm64: dts: qcom: sm8250: Add CPU capacities and energy model")
    Signed-off-by: Vincent Guittot <vincent.guittot@linaro.org>
    Link: https://lore.kernel.org/r/20230615154852.130076-1-vincent.guittot@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2d994433e2d727f373313657810dcf2bba925ae6
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Thu Jun 15 11:34:21 2023 +0300

    Revert "arm64: dts: qcom: msm8996: rename labels for HDMI nodes"

    [ Upstream commit 2b812caf5f64df959555e48dfc7bf8f061d9fe8f ]

    The commit f43b6dc7d56e ("arm64: dts: qcom: msm8996: rename labels for
    HDMI nodes") is broken, it changes all the HDMI node names,
    compatible strings instead of changing just node aliases. Revert the
    commit in order to land a proper clean version.

    Reported-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Fixes: f43b6dc7d56e ("arm64: dts: qcom: msm8996: rename labels for HDMI nodes")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230615083422.350297-2-dmitry.baryshkov@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3e8812eb18769635bf4fec5d8e2a918349df961a
Author: Konrad Dybcio <konrad.dybcio@somainline.org>
Date:   Wed Jun 14 13:35:37 2023 +0200

    arm64: dts: qcom: sm6350: Fix ZAP region

    [ Upstream commit 44bcded2be4fe9b9d0b6e48075c9947b75c0af63 ]

    The previous ZAP region definition was wrong. Fix it.
    Note this is not a device-specific fixup, but a fixup to the generic
    PIL load address.

    Fixes: 5f82b9cda61e ("arm64: dts: qcom: Add SM6350 device tree")
    Signed-off-by: Konrad Dybcio <konrad.dybcio@somainline.org>
```

```
    Reviewed-by: Luca Weiss <luca.weiss@fairphone.com>
    Signed-off-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230315-topic-lagoon_gpu-v2-6-afcdfb18bb13@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6381e6769d8ce3d0734ed88019d0ff74311d9a33
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Mon Jun 12 06:16:23 2023 +0300

    arm64: dts: qcom: sm8150: use proper DSI PHY compatible

    [ Upstream commit 3091e5820a367f3368132f57e0a9ba6d545da15d ]

    The DSI PHY on the Qualcomm SM8150 platform requires platform-specific
    handling. Use the proper SoC-specific compatible string for the DSI
    PHYs.

    Reported-by: Degdag Mohamed <degdagmohamed@gmail.com>
    Fixes: 2ef3bb17c45c ("arm64: dts: qcom: sm8150: Add DISPCC node")
    Cc: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Marijn Suijten <marijn.suijten@somainline.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230612031623.3620155-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c6de94d1a8bd00d3d834910b1b8b0e1327d2f7b2
Author: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
Date:   Mon Jun 19 15:54:04 2023 +0300

    arm64: defconfig: enable Qualcomm MSM8996 Global Clock Controller as built-in

    [ Upstream commit dc015a3a6d6986c41a7bd12fb205a282f685e328 ]

    The commit 8f680c287445 ("arm64: defconfig: Switch msm8996 clk drivers
    to module") switched CONFIG_MSM_MMCC_8996 to module, which also resulted
    in CONFIG_MSM_GCC_8996 being switched to module. This breaks useful
    bootflow for Qualcomm MSM8996 / APQ8096 platforms, because the serial is
    not enabled anymore until the GCC module is loaded.

    Reported-by: Rob Clark <robdclark@gmail.com>
    Fixes: 8f680c287445 ("arm64: defconfig: Switch msm8996 clk drivers to module")
    Signed-off-by: Dmitry Baryshkov <dmitry.baryshkov@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Link: https://lore.kernel.org/r/20230619125404.562137-1-dmitry.baryshkov@linaro.org
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bb275072f5aef1f909a144a02663bd5762164081
Author: Luca Weiss <luca@z3ntu.xyz>
Date:   Wed Jun 14 18:35:47 2023 +0200

    soc: qcom: ocmem: Fix NUM_PORTS & NUM_MACROS macros

    [ Upstream commit a7b484b1c9332a1ee12e8799d62a11ee3f8e0801 ]

    Since we're using these two macros to read a value from a register, we
    need to use the FIELD_GET instead of the FIELD_PREP macro, otherwise
    we're getting wrong values.

    So instead of:

      [    3.111779] ocmem fdd00000.sram: 2 ports, 1 regions, 512 macros, not interleaved
```

```
    we now get the correct value of:

      [    3.129672] ocmem fdd00000.sram: 2 ports, 1 regions, 2 macros, not interleaved

    Fixes: 88c1e9404f1d ("soc: qcom: add OCMEM driver")
    Reviewed-by: Caleb Connolly <caleb.connolly@linaro.org>
    Reviewed-by: Konrad Dybcio <konrad.dybcio@linaro.org>
    Signed-off-by: Luca Weiss <luca@z3ntu.xyz>
    Link: https://lore.kernel.org/r/20230506-msm8226-ocmem-v3-1-79da95a2581f@z3ntu.xyz
    Signed-off-by: Bjorn Andersson <andersson@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9801597014a4937b5ce5474f82ab64ddd69add34
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Fri Jun 30 21:48:36 2023 -0700

    ASoC: stac9766: fix build errors with REGMAP_AC97

    [ Upstream commit c70064b96f509daa78f57992aeabcf274fb2fed4 ]

    Select REGMAP_AC97 to fix these build errors:

    ERROR: modpost: "regmap_ac97_default_volatile" [sound/soc/codecs/snd-soc-stac9766.ko]
undefined!
    ERROR: modpost: "__regmap_init_ac97" [sound/soc/codecs/snd-soc-stac9766.ko] undefined!

    Fixes: 6bbf787bb70c ("ASoC: stac9766: Convert to regmap")
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Cc: Lars-Peter Clausen <lars@metafoo.de>
    Cc: Mark Brown <broonie@kernel.org>
    Cc: Liam Girdwood <lgirdwood@gmail.com>
    Cc: alsa-devel@alsa-project.org
    Link: https://lore.kernel.org/r/20230701044836.18789-1-rdunlap@infradead.org
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 85b969c917265e07fb71465de6367e100d18365d
Author: Sui Jingfeng <suijingfeng@loongson.cn>
Date:   Sun Jul 9 18:05:14 2023 +0800

    drm/hyperv: Fix a compilation issue because of not including screen_info.h

    [ Upstream commit 8d1077cf2e43b15fefd76ebec2b71541eb27ef2c ]

    Fixes the following build errors on arm64:

    drivers/video/fbdev/hyperv_fb.c: In function 'hvfb_getmem':
    >> drivers/video/fbdev/hyperv_fb.c:1033:24: error: 'screen_info' undeclared (first use
in this function)
        1033 |                 base = screen_info.lfb_base;
             |                        ^~~~~~~~~~~
    drivers/video/fbdev/hyperv_fb.c:1033:24: note: each undeclared identifier is reported
only once for each function it appears in

    >> drivers/gpu/drm/hyperv/hyperv_drm_drv.c:75:54: error: 'screen_info' undeclared (first
use in this function)
          75 |         drm_aperture_remove_conflicting_framebuffers(screen_info.lfb_base,
             |                                                      ^~~~~~~~~~~
    drivers/gpu/drm/hyperv/hyperv_drm_drv.c:75:54: note: each undeclared identifier is
reported only once for each function it appears in

    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/oe-kbuild-all/202307090823.nxnT8Kk5-lkp@intel.com/
    Fixes: 81d2393485f0 ("fbdev/hyperv-fb: Do not set struct fb_info.apertures")
```

```
    Fixes: 8b0d13545b09 ("efi: Do not include <linux/screen_info.h> from EFI header")
    Signed-off-by: Sui Jingfeng <suijingfeng@loongson.cn>
    Reviewed-by: Thomas Zimmermann <tzimmermann@suse.de>
    Signed-off-by: Thomas Zimmermann <tzimmermann@suse.de>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230709100514.703759-1-
suijingfeng@loongson.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 999ddaaad9774bdfd99eb8eaf474b88f9190418f
Author: Wesley Chalmers <Wesley.Chalmers@amd.com>
Date:   Thu Nov 3 22:29:31 2022 -0400

    drm/amd/display: Do not set drr on pipe commit

    [ Upstream commit 09c8cbedba5fa85f15ac91ed74848aceff69f8e5 ]

    [WHY]
    Writing to DRR registers such as OTG_V_TOTAL_MIN on the same frame as a
    pipe commit can cause underflow.

    [HOW]
    Move DMUB p-state delegate into optimze_bandwidth; enabling FAMS sets
    optimized_required.

    This change expects that Freesync requests are blocked when
    optimized_required is true.

    Fixes: 613a7956deb3 ("drm/amd/display: Add monitor specific edid quirk")
    Reviewed-by: Rodrigo Siqueira <Rodrigo.Siqueira@amd.com>
    Signed-off-by: Wesley Chalmers <Wesley.Chalmers@amd.com>
    Tested-by: Daniel Wheeler <daniel.wheeler@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Acked-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 31bed65eecbc5ce57592cfe31947eaa64e3d678e
Author: Baokun Li <libaokun1@huawei.com>
Date:   Fri Jun 30 19:08:21 2023 +0800

    quota: fix dqput() to follow the guarantees dquot_srcu should provide

    [ Upstream commit dabc8b20756601b9e1cc85a81d47d3f98ed4d13a ]

    The dquot_mark_dquot_dirty() using dquot references from the inode
    should be protected by dquot_srcu. quota_off code takes care to call
    synchronize_srcu(&dquot_srcu) to not drop dquot references while they
    are used by other users. But dquot_transfer() breaks this assumption.
    We call dquot_transfer() to drop the last reference of dquot and add
    it to free_dquots, but there may still be other users using the dquot
    at this time, as shown in the function graph below:

           cpu1                 cpu2
    ____________________|____________________
    wb_do_writeback         CHOWN(1)
     ...
      ext4_da_update_reserve_space
       dquot_claim_block
        ...
         dquot_mark_dquot_dirty // try to dirty old quota
          test_bit(DQ_ACTIVE_B, &dquot->dq_flags) // still ACTIVE
          if (test_bit(DQ_MOD_B, &dquot->dq_flags))
          // test no dirty, wait dq_list_lock
                        ...
                         dquot_transfer
```

```
                            __dquot_transfer
                            dqput_all(transfer_from) // rls old dquot
                             dqput // last dqput
                              dquot_release
                               clear_bit(DQ_ACTIVE_B, &dquot->dq_flags)
                              atomic_dec(&dquot->dq_count)
                              put_dquot_last(dquot)
                               list_add_tail(&dquot->dq_free, &free_dquots)
                               // add the dquot to free_dquots
          if (!test_and_set_bit(DQ_MOD_B, &dquot->dq_flags))
            add dqi_dirty_list // add released dquot to dirty_list
```

This can cause various issues, such as dquot being destroyed by dqcache_shrink_scan() after being added to free_dquots, which can trigger a UAF in dquot_mark_dquot_dirty(); or after dquot is added to free_dquots and then to dirty_list, it is added to free_dquots again after dquot_writeback_dquots() is executed, which causes the free_dquots list to be corrupted and triggers a UAF when dqcache_shrink_scan() is called for freeing dquot twice.

As Honza said, we need to fix dquot_transfer() to follow the guarantees dquot_srcu should provide. But calling synchronize_srcu() directly from dquot_transfer() is too expensive (and mostly unnecessary). So we add dquot whose last reference should be dropped to the new global dquot list releasing_dquots, and then queue work item which would call synchronize_srcu() and after that perform the final cleanup of all the dquots on releasing_dquots.

Fixes: 4580b30ea887 ("quota: Do not dirty bad dquots")
Suggested-by: Jan Kara <jack@suse.cz>
Signed-off-by: Baokun Li <libaokun1@huawei.com>
Signed-off-by: Jan Kara <jack@suse.cz>
Message-Id: <20230630110822.3881712-5-libaokun1@huawei.com>
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a71166a9a2abb01c10ea05f02987919ddad64cbd
Author: Baokun Li <libaokun1@huawei.com>
Date: Fri Jun 30 19:08:20 2023 +0800

quota: add new helper dquot_active()

[Upstream commit 33bcfafc48cb186bc4bbcea247feaa396594229e]

Add new helper function dquot_active() to make the code more concise.

Signed-off-by: Baokun Li <libaokun1@huawei.com>
Signed-off-by: Jan Kara <jack@suse.cz>
Message-Id: <20230630110822.3881712-4-libaokun1@huawei.com>
Stable-dep-of: dabc8b207566 ("quota: fix dqput() to follow the guarantees dquot_srcu should provide")
Signed-off-by: Sasha Levin <sashal@kernel.org>

commit de2ad5a45c057e99543b3b9548731e4b37586243
Author: Baokun Li <libaokun1@huawei.com>
Date: Fri Jun 30 19:08:19 2023 +0800

quota: rename dquot_active() to inode_quota_active()

[Upstream commit 4b9bdfa16535de8f49bf954aeed0f525ee2fc322]

Now we have a helper function dquot_dirty() to determine if dquot has DQ_MOD_B bit. dquot_active() can easily be misunderstood as a helper function to determine if dquot has DQ_ACTIVE_B bit. So we avoid this by renaming it to inode_quota_active() and later on we will add the helper function dquot_active() to determine if dquot has DQ_ACTIVE_B bit.

```
    Signed-off-by: Baokun Li <libaokun1@huawei.com>
    Signed-off-by: Jan Kara <jack@suse.cz>
    Message-Id: <20230630110822.3881712-3-libaokun1@huawei.com>
    Stable-dep-of: dabc8b207566 ("quota: fix dqput() to follow the guarantees dquot_srcu
should provide")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dbcd5283ece91c1edef2e75175b34d5f5e55ce93
Author: Baokun Li <libaokun1@huawei.com>
Date:   Fri Jun 30 19:08:18 2023 +0800

    quota: factor out dquot_write_dquot()

    [ Upstream commit 024128477809f8073d870307c8157b8826ebfd08 ]

    Refactor out dquot_write_dquot() to reduce duplicate code.

    Signed-off-by: Baokun Li <libaokun1@huawei.com>
    Signed-off-by: Jan Kara <jack@suse.cz>
    Message-Id: <20230630110822.3881712-2-libaokun1@huawei.com>
    Stable-dep-of: dabc8b207566 ("quota: fix dqput() to follow the guarantees dquot_srcu
should provide")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 889299e7bbdda75d6ccd0a93e2218b65d527ac23
Author: Ondrej Jirman <megi@xff.cz>
Date:   Sun Jun 18 00:48:25 2023 +0200

    drm: bridge: dw-mipi-dsi: Fix enable/disable of DSI controller

    [ Upstream commit 05aa61334592adb230749ff465b103ee10e63936 ]

    Before this patch, booting to Linux VT and doing a simple:

      echo 2 > /sys/class/graphics/fb0/blank
      echo 0 > /sys/class/graphics/fb0/blank

    would result in failures to re-enable the panel. Mode set callback is
    called only once during boot in this scenario, while calls to
    enable/disable callbacks are balanced afterwards. The driver doesn't
    work unless userspace calls modeset before enabling the CRTC/connector.

    This patch moves enabling of the DSI host from mode_set into pre_enable
    callback, and removes some old hacks where this bridge driver is
    directly calling into other bridge driver's callbacks.

    pre_enable_prev_first flag is set on the panel's bridge so that panel
    drivers will get their prepare function called between DSI host's
    pre_enable and enable callbacks, so that they get a chance to
    perform panel setup while DSI host is already enabled in command
    mode. Otherwise panel's prepare would be called before DSI host
    is enabled, and any DSI communication used in prepare callback
    would fail.

    With all these changes, the enable/disable sequence is now well
    balanced, and host's and panel's callbacks are called in proper order
    documented in the drm_panel API documentation without needing the old
    hacks. (Mainly that panel->prepare is called when DSI host is ready to
    allow the panel driver to send DSI commands and vice versa during
    disable.)

    Tested on Pinephone Pro. Trace of the callbacks follows.

    Before:
```

```
[    1.253882] dw-mipi-dsi-rockchip ff960000.dsi: mode_set
[    1.290732] panel-himax-hx8394 ff960000.dsi.0: prepare
[    1.475576] dw-mipi-dsi-rockchip ff960000.dsi: enable
[    1.475593] panel-himax-hx8394 ff960000.dsi.0: enable

echo 2 > /sys/class/graphics/fb0/blank

[   13.722799] panel-himax-hx8394 ff960000.dsi.0: disable
[   13.774502] dw-mipi-dsi-rockchip ff960000.dsi: post_disable
[   13.774526] panel-himax-hx8394 ff960000.dsi.0: unprepare

echo 0 > /sys/class/graphics/fb0/blank

[   17.735796] panel-himax-hx8394 ff960000.dsi.0: prepare
[   17.923522] dw-mipi-dsi-rockchip ff960000.dsi: enable
[   17.923540] panel-himax-hx8394 ff960000.dsi.0: enable
[   17.944330] dw-mipi-dsi-rockchip ff960000.dsi: failed to write command FIFO
[   17.944335] panel-himax-hx8394 ff960000.dsi.0: sending command 0xb9 failed: -110
[   17.944340] panel-himax-hx8394 ff960000.dsi.0: Panel init sequence failed: -110

echo 2 > /sys/class/graphics/fb0/blank

[  431.148583] panel-himax-hx8394 ff960000.dsi.0: disable
[  431.169259] dw-mipi-dsi-rockchip ff960000.dsi: failed to write command FIFO
[  431.169268] panel-himax-hx8394 ff960000.dsi.0: Failed to enter sleep mode: -110
[  431.169282] dw-mipi-dsi-rockchip ff960000.dsi: post_disable
[  431.169316] panel-himax-hx8394 ff960000.dsi.0: unprepare
[  431.169357] pclk_mipi_dsi0 already disabled

echo 0 > /sys/class/graphics/fb0/blank

[  432.796851] panel-himax-hx8394 ff960000.dsi.0: prepare
[  432.981537] dw-mipi-dsi-rockchip ff960000.dsi: enable
[  432.981568] panel-himax-hx8394 ff960000.dsi.0: enable
[  433.002290] dw-mipi-dsi-rockchip ff960000.dsi: failed to write command FIFO
[  433.002299] panel-himax-hx8394 ff960000.dsi.0: sending command 0xb9 failed: -110
[  433.002312] panel-himax-hx8394 ff960000.dsi.0: Panel init sequence failed: -110

---------------------------------------------------------------------

After:

[    1.248372] dw-mipi-dsi-rockchip ff960000.dsi: mode_set
[    1.248704] dw-mipi-dsi-rockchip ff960000.dsi: pre_enable
[    1.285377] panel-himax-hx8394 ff960000.dsi.0: prepare
[    1.468392] dw-mipi-dsi-rockchip ff960000.dsi: enable
[    1.468421] panel-himax-hx8394 ff960000.dsi.0: enable

echo 2 > /sys/class/graphics/fb0/blank

[   16.210357] panel-himax-hx8394 ff960000.dsi.0: disable
[   16.261315] dw-mipi-dsi-rockchip ff960000.dsi: post_disable
[   16.261339] panel-himax-hx8394 ff960000.dsi.0: unprepare

echo 0 > /sys/class/graphics/fb0/blank

[   19.161453] dw-mipi-dsi-rockchip ff960000.dsi: pre_enable
[   19.197869] panel-himax-hx8394 ff960000.dsi.0: prepare
[   19.382141] dw-mipi-dsi-rockchip ff960000.dsi: enable
[   19.382158] panel-himax-hx8394 ff960000.dsi.0: enable

       (But depends on functionality intorduced in Linux 6.3, so this patch will
        not build on older kernels when applied to older stable branches.)
```

```
    Fixes: 46fc51546d44 ("drm/bridge/synopsys: Add MIPI DSI host controller bridge")
    Signed-off-by: Ondrej Jirman <megi@xff.cz>
    Reviewed-by: Sam Ravnborg <sam@ravnborg.org>
    Signed-off-by: Robert Foss <rfoss@kernel.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230617224915.1923630-1-megi@xff.cz
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b6e53bc905912325d19da5f43554adfef369832d
Author: Phil Elwell <phil@raspberrypi.com>
Date:   Wed Jun 21 16:32:29 2023 +0100

    ASoC: cs43130: Fix numerator/denominator mixup

    [ Upstream commit a9e7c964cea4fb1541cc81a11d1b2fd135f4cf38 ]

    In converting to using the standard u16_fract type, commit [1] made the
    obvious mistake and failed to take account of the difference in
    numerator and denominator ordering, breaking all uses of the cs43130
    codec.

    Fix it.

    [1] commit e14bd35ef446 ("ASoC: cs43130: Re-use generic struct u16_fract")

    Fixes: e14bd35ef446 ("ASoC: cs43130: Re-use generic struct u16_fract")
    Signed-off-by: Phil Elwell <phil@raspberrypi.com>
    Reviewed-by: Andy Shevchenko <andriy.shevchenko@linux.intel.com>
    Acked-by: Charles Keepax <ckeepax@opensource.cirrus.com>
    Link: https://lore.kernel.org/r/20230621153229.1944132-1-phil@raspberrypi.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ccd338ad38655d186be77c5117139d4c32e4d635
Author: yixuanjiang <yixuanjiang@google.com>
Date:   Mon Jun 19 11:31:27 2023 +0800

    ASoC: soc-compress: Fix deadlock in soc_compr_open_fe

    [ Upstream commit 2222214749a9969e09454b9ba7febfdfb09c1c8d ]

    Modify the error handling flow by release lock.
    The require mutex will keep holding if open fail.

    Fixes: aa9ff6a4955f ("ASoC: soc-compress: Reposition and add pcm_mutex")
    Signed-off-by: yixuanjiang <yixuanjiang@google.com>
    Link: https://lore.kernel.org/r/20230619033127.2522477-1-yixuanjiang@google.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d029fd68b24da2d39c504e7604ed3466eab8d933
Author: Marek Vasut <marex@denx.de>
Date:   Thu Jun 15 17:28:17 2023 +0200

    drm/bridge: tc358764: Fix debug print parameter order

    [ Upstream commit 7f947be02aab5b154427cb5b0fffe858fc387b02 ]

    The debug print parameters were swapped in the output and they were
    printed as decimal values, both the hardware address and the value.
    Update the debug print to print the parameters in correct order, and
    use hexadecimal print for both address and value.

    Fixes: f38b7cca6d0e ("drm/bridge: tc358764: Add DSI to LVDS bridge driver")
    Signed-off-by: Marek Vasut <marex@denx.de>
    Reviewed-by: Robert Foss <rfoss@kernel.org>
```

```
    Signed-off-by: Robert Foss <rfoss@kernel.org>
    Link: https://patchwork.freedesktop.org/patch/msgid/20230615152817.359420-1-
 marex@denx.de
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 2fec6957223bbdd810c1a6dee48ce5b7b1174802
 Author: Kuniyuki Iwashima <kuniyu@amazon.com>
 Date:   Thu Aug 24 09:50:59 2023 -0700

    netrom: Deny concurrent connect().

    [ Upstream commit c2f8fd7949603efb03908e05abbf7726748c8de3 ]

    syzkaller reported null-ptr-deref [0] related to AF_NETROM.
    This is another self-accept issue from the strace log. [1]

    syz-executor creates an AF_NETROM socket and calls connect(), which
    is blocked at that time.  Then, sk->sk_state is TCP_SYN_SENT and
    sock->state is SS_CONNECTING.

      [pid  5059] socket(AF_NETROM, SOCK_SEQPACKET, 0) = 4
      [pid  5059] connect(4, {sa_family=AF_NETROM, sa_data="..." <unfinished ...>

    Another thread calls connect() concurrently, which finally fails
    with -EINVAL.  However, the problem here is the socket state is
    reset even while the first connect() is blocked.

      [pid  5060] connect(4, NULL, 0 <unfinished ...>
      [pid  5060] <... connect resumed>)      = -1 EINVAL (Invalid argument)

    As sk->state is TCP_CLOSE and sock->state is SS_UNCONNECTED, the
    following listen() succeeds.  Then, the first connect() looks up
    itself as a listener and puts skb into the queue with skb->sk itself.
    As a result, the next accept() gets another FD of itself as 3, and
    the first connect() finishes.

      [pid  5060] listen(4, 0 <unfinished ...>
      [pid  5060] <... listen resumed>)       = 0
      [pid  5060] accept(4, NULL, NULL <unfinished ...>
      [pid  5060] <... accept resumed>)       = 3
      [pid  5059] <... connect resumed>)      = 0

    Then, accept4() is called but blocked, which causes the general protection
    fault later.

      [pid  5059] accept4(4, NULL, 0x20000400, SOCK_NONBLOCK <unfinished ...>

    After that, another self-accept occurs by accept() and writev().

      [pid  5060] accept(4, NULL, NULL <unfinished ...>
      [pid  5061] writev(3, [{iov_base=...}] <unfinished ...>
      [pid  5061] <... writev resumed>)       = 99
      [pid  5060] <... accept resumed>)       = 6

    Finally, the leader thread close()s all FDs.  Since the three FDs
    reference the same socket, nr_release() does the cleanup for it
    three times, and the remaining accept4() causes the following fault.

      [pid  5058] close(3)                    = 0
      [pid  5058] close(4)                    = 0
      [pid  5058] close(5)                    = -1 EBADF (Bad file descriptor)
      [pid  5058] close(6)                    = 0
      [pid  5058] <... exit_group resumed>)   = ?
      [   83.456055][ T5059] general protection fault, probably for non-canonical address
 0xdffffc0000000003: 0000 [#1] PREEMPT SMP KASAN
```

```
    To avoid the issue, we need to return an error for connect() if
    another connect() is in progress, as done in __inet_stream_connect().

    [0]:
    general protection fault, probably for non-canonical address 0xdfffc0000000003: 0000
 [#1] PREEMPT SMP KASAN
    KASAN: null-ptr-deref in range [0x0000000000000018-0x000000000000001f]
    CPU: 0 PID: 5059 Comm: syz-executor.0 Not tainted 6.5.0-rc5-syzkaller-00194-
 gace0ab3a4b54 #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
 07/26/2023
    RIP: 0010:__lock_acquire+0x109/0x5de0 kernel/locking/lockdep.c:5012
    Code: 45 85 c9 0f 84 cc 0e 00 00 44 8b 05 11 6e 23 0b 45 85 c0 0f 84 be 0d 00 00 48 ba
 00 00 00 00 00 fc ff df 4c 89 d1 48 c1 e9 03 <80> 3c 11 00 0f 85 e8 40 00 00 49 81 3a a0 69
 48 90 0f 84 96 0d 00
    RSP: 0018:ffffc90003d6f9e0 EFLAGS: 00010006
    RAX: ffff8880244c8000 RBX: 1ffff920007adf6c RCX: 0000000000000003
    RDX: dfffc0000000000 RSI: 0000000000000000 RDI: 0000000000000018
    RBP: 0000000000000001 R08: 0000000000000001 R09: 0000000000000001
    R10: 0000000000000018 R11: 0000000000000000 R12: 0000000000000000
    R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000000
    FS:  00007f51d519a6c0(0000) GS:ffff8880b9800000(0000) knlGS:0000000000000000
    CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
    CR2: 00007f51d5158d58 CR3: 000000002943f000 CR4: 00000000003506f0
    DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
    DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400
    Call Trace:
     <TASK>
     lock_acquire kernel/locking/lockdep.c:5761 [inline]
     lock_acquire+0x1ae/0x510 kernel/locking/lockdep.c:5726
     __raw_spin_lock_irqsave include/linux/spinlock_api_smp.h:110 [inline]
     _raw_spin_lock_irqsave+0x3a/0x50 kernel/locking/spinlock.c:162
     prepare_to_wait+0x47/0x380 kernel/sched/wait.c:269
     nr_accept+0x20d/0x650 net/netrom/af_netrom.c:798
     do_accept+0x3a6/0x570 net/socket.c:1872
     __sys_accept4_file net/socket.c:1913 [inline]
     __sys_accept4+0x99/0x120 net/socket.c:1943
     __do_sys_accept4 net/socket.c:1954 [inline]
     __se_sys_accept4 net/socket.c:1951 [inline]
     __x64_sys_accept4+0x96/0x100 net/socket.c:1951
     do_syscall_x64 arch/x86/entry/common.c:50 [inline]
     do_syscall_64+0x38/0xb0 arch/x86/entry/common.c:80
     entry_SYSCALL_64_after_hwframe+0x63/0xcd
    RIP: 0033:0x7f51d447cae9
    Code: 28 00 00 00 75 05 48 83 c4 28 c3 e8 e1 20 00 00 90 48 89 f8 48 89 f7 48 89 d6 48
 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 b0 ff ff
 ff f7 d8 64 89 01 48
    RSP: 002b:00007f51d519a0c8 EFLAGS: 00000246 ORIG_RAX: 0000000000000120
    RAX: ffffffffffffffda RBX: 00007f51d459bf80 RCX: 00007f51d447cae9
    RDX: 0000000020000400 RSI: 0000000000000000 RDI: 0000000000000004
    RBP: 00007f51d44c847a R08: 0000000000000000 R09: 0000000000000000
    R10: 0000000000000800 R11: 0000000000000246 R12: 0000000000000000
    R13: 000000000000000b R14: 00007f51d459bf80 R15: 00007ffc25c34e48
     </TASK>

    Link: https://syzkaller.appspot.com/text?tag=CrashLog&x=152cdb63a80000 [1]
    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reported-by: syzbot+666c97e4686410e79649@syzkaller.appspotmail.com
    Closes: https://syzkaller.appspot.com/bug?extid=666c97e4686410e79649
    Signed-off-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit b1155098f04a108dbc8069733aefd4f4b1d8b4a6
```

```
Author: Shannon Nelson <shannon.nelson@amd.com>
Date:   Thu Aug 24 09:17:54 2023 -0700

    pds_core: pass opcode to devcmd_wait

    [ Upstream commit 0ea064e74bc8f915aba3f2d0fb3418247a09b73d ]

    Don't rely on the PCI memory for the devcmd opcode because we
    read a 0xff value if the PCI bus is broken, which can cause us
    to report a bogus dev_cmd opcode later.

    Fixes: 523847df1b37 ("pds_core: add devcmd device interfaces")
    Signed-off-by: Shannon Nelson <shannon.nelson@amd.com>
    Reviewed-by: Brett Creeley <brett.creeley@amd.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824161754.34264-6-shannon.nelson@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 410b24635fb992324ac307f8197a50e90dddb5f9
Author: Shannon Nelson <shannon.nelson@amd.com>
Date:   Thu Aug 24 09:17:53 2023 -0700

    pds_core: check for work queue before use

    [ Upstream commit 969cfd4c8ca50c32901342cdd3d677c3ffe61371 ]

    Add a check that the wq exists before queuing up work for a
    failed devcmd, as the PF is responsible for health and the VF
    doesn't have a wq.

    Fixes: c2dbb0904310 ("pds_core: health timer and workqueue")
    Signed-off-by: Shannon Nelson <shannon.nelson@amd.com>
    Reviewed-by: Brett Creeley <brett.creeley@amd.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824161754.34264-5-shannon.nelson@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d4fb4628740d89768a6b97c5b45d657c3cbd184
Author: Shannon Nelson <shannon.nelson@amd.com>
Date:   Thu Aug 24 09:17:52 2023 -0700

    pds_core: no reset command for VF

    [ Upstream commit 95e383226d6fcda6c217912f11edf8d74de9cc85 ]

    The VF doesn't need to send a reset command, and in a PCI reset
    scenario it might not have a valid IO space to write to anyway.

    Fixes: 523847df1b37 ("pds_core: add devcmd device interfaces")
    Signed-off-by: Shannon Nelson <shannon.nelson@amd.com>
    Reviewed-by: Brett Creeley <brett.creeley@amd.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824161754.34264-4-shannon.nelson@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit faefd384b4813284d4485d92a2649dd39648e04b
Author: Shannon Nelson <shannon.nelson@amd.com>
Date:   Thu Aug 24 09:17:51 2023 -0700

    pds_core: no health reporter in VF

    [ Upstream commit e48b894a1db7f6ce66bff0402ab21ff9f0e56034 ]
```

```
    Make sure the health reporter is set up before we use it in
    our devlink health updates, especially since the VF doesn't
    set up the health reporter.

    Fixes: 25b450c05a49 ("pds_core: add devlink health facilities")
    Signed-off-by: Shannon Nelson <shannon.nelson@amd.com>
    Reviewed-by: Brett Creeley <brett.creeley@amd.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824161754.34264-3-shannon.nelson@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b84cccf13c156f430944ed892b57b1184410470e
Author: Shannon Nelson <shannon.nelson@amd.com>
Date:   Thu Aug 24 09:17:50 2023 -0700

    pds_core: protect devlink callbacks from fw_down state

    [ Upstream commit 91202ce78fcd070982a115f0bf6f328af619aa00 ]

    Don't access structs that have been cleared when in the fw_down
    state and the various structs have been cleaned and are waiting
    to recover.  This caused a panic on rmmod when already in fw_down
    and devlink_param_unregister() tried to check the parameters.

    Fixes: 40ced8944536 ("pds_core: devlink params for enabling VIF support")
    Signed-off-by: Shannon Nelson <shannon.nelson@amd.com>
    Reviewed-by: Brett Creeley <brett.creeley@amd.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824161754.34264-2-shannon.nelson@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit eb07894c51c7d6bb8d00948a3e6e7b52c791e93e
Author: Budimir Markovic <markovicbudimir@gmail.com>
Date:   Thu Aug 24 01:49:05 2023 -0700

    net/sched: sch_hfsc: Ensure inner classes have fsc curve

    [ Upstream commit b3d26c5702c7d6c45456326e56d2ccf3f103e60f ]

    HFSC assumes that inner classes have an fsc curve, but it is currently
    possible for classes without an fsc curve to become parents. This leads
    to bugs including a use-after-free.

    Don't allow non-root classes without HFSC_FSC to become parents.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Reported-by: Budimir Markovic <markovicbudimir@gmail.com>
    Signed-off-by: Budimir Markovic <markovicbudimir@gmail.com>
    Acked-by: Jamal Hadi Salim <jhs@mojatatu.com>
    Link: https://lore.kernel.org/r/20230824084905.422-1-markovicbudimir@gmail.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 99b316bb089ec221d69e88ddee8145371670988c
Author: Alex Austin <alex.austin@amd.com>
Date:   Thu Aug 24 17:46:57 2023 +0100

    sfc: Check firmware supports Ethernet PTP filter

    [ Upstream commit c4413a20fa6d7c4888009fb7dd391685f196cd36 ]

    Not all firmware variants support RSS filters. Do not fail all PTP
```

```
    functionality when raw ethernet PTP filters fail to insert.

    Fixes: e4616f64726b ("sfc: support PTP over Ethernet")
    Signed-off-by: Alex Austin <alex.austin@amd.com>
    Acked-by: Edward Cree <ecree.xilinx@gmail.com>
    Reviewed-by: Pieter Jansen van Vuuren <pieter.jansen-van-vuuren@amd.com>
    Link: https://lore.kernel.org/r/20230824164657.42379-1-alex.austin@amd.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b3dc4c0d4a122570b9746191411b48c4803c3c84
Author: Suman Ghosh <sumang@marvell.com>
Date:   Thu Aug 24 13:40:32 2023 +0530

    cteonxt2-pf: Fix backpressure config for multiple PFC priorities to work simultaneously

    [ Upstream commit 597d0ec0e4ca6a912affea4cc94df08959e9ec74 ]

    MAC (CGX or RPM) asserts backpressure at TL3 or TL2 node of the egress
    hierarchical scheduler tree depending on link level config done. If
    there are multiple PFC priorities enabled at a time and for all such
    flows to backoff, each priority will have to assert backpressure at
    different TL3/TL2 scheduler nodes and these flows will need to submit
    egress pkts to these nodes.

    Current PFC configuration has an issue where in only one backpressure
    scheduler node is being allocated which is resulting in only one PFC
    priority to work. This patch fixes this issue.

    Fixes: 99c969a83d82 ("octeontx2-pf: Add egress PFC support")
    Signed-off-by: Suman Ghosh <sumang@marvell.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824081032.436432-4-sumang@marvell.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8ce75ae99761887c9ea15f81a790d4112a334d6d
Author: Hariprasad Kelam <hkelam@marvell.com>
Date:   Thu Aug 24 13:40:31 2023 +0530

    octeontx2-af: CN10KB: fix PFC configuration

    [ Upstream commit 47bcc9c1cf6aa60156c7532983090e86d9d171b6 ]

    Suppose user has enabled pfc with prio 0,1 on a PF netdev(eth0)
            dcb pfc set dev eth0 prio-pfc o:on 1:on
    later user enabled pfc priorities 2 and 3 on the VF interface(eth1)
            dcb pfc set dev eth1 prio-pfc 2:on 3:on

    Instead of enabling pfc on all priorities (0..3), the driver only
    enables on priorities 2,3. This patch corrects the issue by using
    the proper CSR address.

    Fixes: b9d0fedc6234 ("octeontx2-af: cn10kb: Add RPM_USX MAC support")
    Signed-off-by: Hariprasad Kelam <hkelam@marvell.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824081032.436432-3-sumang@marvell.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f72dec7d5ad9387545a8615fa74397bc9e3e5cc1
Author: Suman Ghosh <sumang@marvell.com>
Date:   Thu Aug 24 13:40:30 2023 +0530

    octeontx2-pf: Fix PFC TX scheduler free
```

```
    [ Upstream commit a9ac2e18779597f280d68a5b5f5bdd51a34080fa ]

    During PFC TX schedulers free, flag TXSCHQ_FREE_ALL was being set
    which caused free up all schedulers other than the PFC schedulers.
    This patch fixes that to free only the PFC Tx schedulers.

    Fixes: 99c969a83d82 ("octeontx2-pf: Add egress PFC support")
    Signed-off-by: Suman Ghosh <sumang@marvell.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230824081032.436432-2-sumang@marvell.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 64a575b76c19d9848d81d38de72314c4439dd6b8
Author: Biju Das <biju.das.jz@bp.renesas.com>
Date:   Thu Aug 24 21:44:54 2023 +0100

    hwmon: (tmp513) Fix the channel number in tmp51x_is_visible()

    [ Upstream commit d103337e38e7e64c3d915029e947b1cb0b512737 ]

    The supported channels for this driver are {0..3}. Fix the incorrect
    channel in tmp51x_is_visible().

    Reported-by: Guenter Roeck <linux@roeck-us.net>
    Closes: https://lore.kernel.org/all/ea0eccc0-a29f-41e4-9049-a1a13f8b16f1@roeck-us.net/
    Fixes: 59dfa75e5d82 ("hwmon: Add driver for Texas Instruments TMP512/513 sensor chips.")
    Signed-off-by: Biju Das <biju.das.jz@bp.renesas.com>
    Link: https://lore.kernel.org/r/20230824204456.401580-2-biju.das.jz@bp.renesas.com
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3b0a97365c2c5a4f1d880f6def31f87c4fc78b32
Author: Adam Guerin <adam.guerin@intel.com>
Date:   Mon Aug 14 16:52:30 2023 +0100

    crypto: qat - fix crypto capability detection for 4xxx

    [ Upstream commit fab9516f02b418e37d3cde6c21c316085262aece ]

    When extending the capability detection logic for 4xxx devices the
    SMx algorithms were accidentally missed.
    Enable these SMx capabilities by default for QAT GEN4 devices.

    Check for device variants where the SMx algorithms are explicitly
    disabled by the GEN4 hardware. This is indicated in fusectl1
    register.
    Mask out SM3 and SM4 based on a bit specific to those algorithms.
    Mask out SM2 if the PKE slice is not present.

    Fixes: 4b44d28c715d ("crypto: qat - extend crypto capability detection for 4xxx")
    Signed-off-by: Adam Guerin <adam.guerin@intel.com>
    Reviewed-by: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
    Reviewed-by: Fiona Trahe <fiona.trahe@intel.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2f7c3ea908ce0830ec5d7dd85ad11c8ad84770f7
Author: Vadim Pasternak <vadimp@nvidia.com>
Date:   Thu Aug 24 15:43:10 2023 +0200

    mlxsw: core_hwmon: Adjust module label names based on MTCAP sensor counter

    [ Upstream commit 3fc134a07438055fc93ce1bbacf2702ddd09500c ]
```

```
    Transceiver module temperature sensors are indexed after ASIC and
    platform sensors. The current label printing method does not take this
    into account and simply prints the index of the transceiver module
    sensor.

    On new systems that have platform sensors this results in incorrect
    (shifted) transceiver module labels being printed:

    $ sensors
    [...]
    front panel 002:  +37.0°C  (crit = +70.0°C, emerg = +75.0°C)
    front panel 003:  +47.0°C  (crit = +70.0°C, emerg = +75.0°C)
    [...]

    Fix by taking the sensor count into account. After the fix:

    $ sensors
    [...]
    front panel 001:  +37.0°C  (crit = +70.0°C, emerg = +75.0°C)
    front panel 002:  +47.0°C  (crit = +70.0°C, emerg = +75.0°C)
    [...]

    Fixes: a53779de6a0e ("mlxsw: core: Add QSFP module temperature label attribute to
hwmon")
    Signed-off-by: Vadim Pasternak <vadimp@nvidia.com>
    Reviewed-by: Ido Schimmel <idosch@nvidia.com>
    Signed-off-by: Petr Machata <petrm@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e137f1b7670b28bd1413eccb7b12b759f838a029
Author: Vadim Pasternak <vadimp@nvidia.com>
Date:   Thu Aug 24 15:43:09 2023 +0200

    mlxsw: i2c: Limit single transaction buffer size

    [ Upstream commit d7248f1cc835bd80e936dc5b2d94b149bdd0077d ]

    Maximum size of buffer is obtained from underlying I2C adapter and in
    case adapter allows I2C transaction buffer size greater than 100 bytes,
    transaction will fail due to firmware limitation.

    As a result driver will fail initialization.

    Limit the maximum size of transaction buffer by 100 bytes to fit to
    firmware.

    Remove unnecessary calculation:
    max_t(u16, MLXSW_I2C_BLK_DEF, quirk_size).
    This condition can not happened.

    Fixes: 3029a693beda ("mlxsw: i2c: Allow flexible setting of I2C transactions size")
    Signed-off-by: Vadim Pasternak <vadimp@nvidia.com>
    Reviewed-by: Petr Machata <petrm@nvidia.com>
    Signed-off-by: Petr Machata <petrm@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 92581d37bdaaafb1e9d0bf4aec5603e3c4c766d3
Author: Vadim Pasternak <vadimp@nvidia.com>
Date:   Thu Aug 24 15:43:08 2023 +0200

    mlxsw: i2c: Fix chunk size setting in output mailbox buffer
```

```
    [ Upstream commit 146c7c330507c0384bf29d567186632bfe975927 ]

    The driver reads commands output from the output mailbox. If the size
    of the output mailbox is not a multiple of the transaction /
    block size, then the driver will not issue enough read transactions
    to read the entire output, which can result in driver initialization
    errors.

    Fix by determining the number of transactions using DIV_ROUND_UP().

    Fixes: 3029a693beda ("mlxsw: i2c: Allow flexible setting of I2C transactions size")
    Signed-off-by: Vadim Pasternak <vadimp@nvidia.com>
    Reviewed-by: Ido Schimmel <idosch@nvidia.com>
    Signed-off-by: Petr Machata <petrm@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c5b84d1ac88ce8f89f3ff9c850fc6f833f1385f6
Author: Christophe Leroy <christophe.leroy@csgroup.eu>
Date:   Wed Aug 23 15:21:43 2023 +0200

    kunit: Fix checksum tests on big endian CPUs

    [ Upstream commit b38460bc463c54e0c15ff3b37e81f7e2059bb9bb ]

    On powerpc64le checksum kunit tests work:

    [    2.011457][    T1]     KTAP version 1
    [    2.011662][    T1]     # Subtest: checksum
    [    2.011848][    T1]     1..3
    [    2.034710][    T1]     ok 1 test_csum_fixed_random_inputs
    [    2.079325][    T1]     ok 2 test_csum_all_carry_inputs
    [    2.127102][    T1]     ok 3 test_csum_no_carry_inputs
    [    2.127202][    T1] # checksum: pass:3 fail:0 skip:0 total:3
    [    2.127533][    T1] # Totals: pass:3 fail:0 skip:0 total:3
    [    2.127956][    T1] ok 1 checksum

    But on powerpc64 and powerpc32 they fail:

    [    1.859890][    T1]     KTAP version 1
    [    1.860041][    T1]     # Subtest: checksum
    [    1.860201][    T1]     1..3
    [    1.861927][   T58]     # test_csum_fixed_random_inputs: ASSERTION FAILED at
lib/checksum_kunit.c:243
    [    1.861927][   T58]     Expected result == expec, but
    [    1.861927][   T58]         result == 54991 (0xd6cf)
    [    1.861927][   T58]         expec == 33316 (0x8224)
    [    1.863742][    T1]     not ok 1 test_csum_fixed_random_inputs
    [    1.864520][   T60]     # test_csum_all_carry_inputs: ASSERTION FAILED at
lib/checksum_kunit.c:267
    [    1.864520][   T60]     Expected result == expec, but
    [    1.864520][   T60]         result == 255 (0xff)
    [    1.864520][   T60]         expec == 65280 (0xff00)
    [    1.868820][    T1]     not ok 2 test_csum_all_carry_inputs
    [    1.869977][   T62]     # test_csum_no_carry_inputs: ASSERTION FAILED at
lib/checksum_kunit.c:306
    [    1.869977][   T62]     Expected result == expec, but
    [    1.869977][   T62]         result == 64515 (0xfc03)
    [    1.869977][   T62]         expec == 0 (0x0)
    [    1.872060][    T1]     not ok 3 test_csum_no_carry_inputs
    [    1.872102][    T1] # checksum: pass:0 fail:3 skip:0 total:3
    [    1.872458][    T1] # Totals: pass:0 fail:3 skip:0 total:3
    [    1.872791][    T1] not ok 3 checksum

    This is because all expected values were calculated for X86 which
```

```
    is little endian. On big endian systems all precalculated 16 bits
    halves must be byte swapped.

    And this is confirmed by a huge amount of sparse errors when building
    with C=2

    So fix all sparse errors and it will naturally work on all endianness.

    Fixes: 688eb8191b47 ("x86/csum: Improve performance of `csum_partial`")
    Signed-off-by: Christophe Leroy <christophe.leroy@csgroup.eu>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bd386430bca007298fdeffdd705a0e20037f7ecf
Author: Jinjie Ruan <ruanjinjie@huawei.com>
Date:   Thu Aug 24 14:43:36 2023 +0800

    net: arcnet: Do not call kfree_skb() under local_irq_disable()

    [ Upstream commit 786c96e92fb9e854cb8b0cb7399bb2fb28e15c4b ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with hardware interrupts being disabled.
    So replace kfree_skb() with dev_kfree_skb_irq() under
    local_irq_disable(). Compile tested only.

    Fixes: 05fcd31cc472 ("arcnet: add err_skb package for package status feedback")
    Signed-off-by: Jinjie Ruan <ruanjinjie@huawei.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 78d08e77554dba1ff3540c8359b2a715e2b3cbbe
Author: Ratheesh Kannoth <rkannoth@marvell.com>
Date:   Thu Aug 24 08:33:01 2023 +0530

    octeontx2-pf: fix page_pool creation fail for rings > 32k

    [ Upstream commit 49fa4b0d06705a24a81bb8be6eb175059b77f0a7 ]

    octeontx2 driver calls page_pool_create() during driver probe()
    and fails if queue size > 32k. Page pool infra uses these buffers
    as shock absorbers for burst traffic. These pages are pinned down
    over time as working sets varies, due to the recycling nature
    of page pool, given page pool (currently) don't have a shrinker
    mechanism, the pages remain pinned down in ptr_ring.
    Instead of clamping page_pool size to 32k at
    most, limit it even more to 2k to avoid wasting memory.

    This have been tested on octeontx2 CN10KA hardware.
    TCP and UDP tests using iperf shows no performance regressions.

    Fixes: b2e3406a38f0 ("octeontx2-pf: Add support for page pool")
    Suggested-by: Alexander Lobakin <aleksander.lobakin@intel.com>
    Reviewed-by: Sunil Goutham <sgoutham@marvell.com>
    Signed-off-by: Ratheesh Kannoth <rkannoth@marvell.com>
    Acked-by: Jesper Dangaard Brouer <hawk@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cce9dd4498e85d4b48f77ff4d1d40d6f2128f20d
Author: Jacob Keller <jacob.e.keller@intel.com>
Date:   Wed Aug 23 08:18:14 2023 -0700

    ice: avoid executing commands on other ports when driving sync
```

```
    [ Upstream commit 0aacec49c29e7c5b1487e859b0c0a42388c34092 ]

    The ice hardware has a synchronization mechanism used to drive the
    simultaneous application of commands on both PHY ports and the source timer
    in the MAC.

    When issuing a sync via ice_ptp_exec_tmr_cmd(), the hardware will
    simultaneously apply the commands programmed for the main timer and each
    PHY port. Neither the main timer command register, nor the PHY port command
    registers auto clear on command execution.

    During the execution of a timer command intended for a single port on E822
    devices, such as those used to configure a PHY during link up, the driver
    is not correctly clearing the previous commands.

    This results in unintentionally executing the last programmed command on
    the main timer and other PHY ports whenever performing reconfiguration on
    E822 ports after link up. This results in unintended side effects on other
    timers, depending on what command was previously programmed.

    To fix this, the driver must ensure that the main timer and all other PHY
    ports are properly initialized to perform no action.

    The enumeration for timer commands does not include an enumeration value
    for doing nothing. Introduce ICE_PTP_NOP for this purpose. When writing a
    timer command to hardware, leave the command bits set to zero which
    indicates that no operation should be performed on that port.

    Modify ice_ptp_one_port_cmd() to always initialize all ports. For all ports
    other than the one being configured, write their timer command register to
    ICE_PTP_NOP. This ensures that no side effect happens on the timer command.

    To fix this for the PHY ports, modify ice_ptp_one_port_cmd() to always
    initialize all other ports to ICE_PTP_NOP. This ensures that no side
    effects happen on the other ports.

    Call ice_ptp_src_cmd() with a command value if ICE_PTP_NOP in
    ice_sync_phy_timer_e822() and ice_start_phy_timer_e822().

    With both of these changes, the driver should no longer execute a stale
    command on the main timer or another PHY port when reconfiguring one of the
    PHY ports after link up.

    Fixes: 3a7496234d17 ("ice: implement basic E822 PTP support")
    Signed-off-by: Siddaraju DH <siddaraju.dh@intel.com>
    Signed-off-by: Jacob Keller <jacob.e.keller@intel.com>
    Tested-by: Sunitha Mekala <sunithax.d.mekala@intel.com> (A Contingent worker at Intel)
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5736bf7341d7781a9ed65c29adcfdb6efafb1c72
Author: Wang Ming <machel@vivo.com>
Date:   Thu Jul 13 11:03:44 2023 +0800

    wifi: ath9k: use IS_ERR() with debugfs_create_dir()

    [ Upstream commit 1e4134610d93271535ecf900a676e1f094e9944c ]

    The debugfs_create_dir() function returns error pointers,
    it never returns NULL. Most incorrect error checks were fixed,
    but the one in ath9k_htc_init_debug() was forgotten.

    Fix the remaining error check.
```

```
    Fixes: e5facc75fa91 ("ath9k_htc: Cleanup HTC debugfs")
    Signed-off-by: Wang Ming <machel@vivo.com>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230713030358.12379-1-machel@vivo.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d7711190906ae06a8acec8f3df425cd7a5efd743
Author: Rahul Rameshbabu <rrameshbabu@nvidia.com>
Date:   Mon Aug 21 16:05:54 2023 -0700

    net/mlx5: Dynamic cyclecounter shift calculation for PTP free running clock

    [ Upstream commit 84a58e60038fa0366006977dba85eae16b2e3d78 ]

    Use a dynamic calculation to determine the shift value for the internal
    timer cyclecounter that will lead to the highest precision frequency
    adjustments. Previously used a constant for the shift value assuming all
    devices supported by the driver had a nominal frequency of 1GHz. However,
    there are devices that operate at different frequencies. The previous shift
    value constant would break the PHC functionality for those devices.

    Reported-by: Vadim Fedorenko <vadim.fedorenko@linux.dev>
    Closes: https://lore.kernel.org/netdev/20230815151507.3028503-1-vadfed@meta.com/
    Fixes: 6a4010927562 ("net/mlx5: Update cyclecounter shift value to improve ptp free
running mode precision")
    Signed-off-by: Rahul Rameshbabu <rrameshbabu@nvidia.com>
    Tested-by: Vadim Fedorenko <vadim.fedorenko@linux.dev>
    Reviewed-by: Jacob Keller <jacob.e.keller@intel.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Acked-by: Saeed Mahameed <saeedm@nvidia.com>
    Link: https://lore.kernel.org/r/20230821230554.236210-1-rrameshbabu@nvidia.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b343c2717bb27b7596131f3ffaa90ebe232738a9
Author: Qi Zheng <zhengqi.arch@bytedance.com>
Date:   Thu Aug 10 09:32:41 2023 +0000

    arm64: mm: use ptep_clear() instead of pte_clear() in clear_flush()

    [ Upstream commit 00de2c9f26b15f1a6f2af516dd8ec5f8d28189b7 ]

    In clear_flush(), the original pte may be a present entry, so we should
    use ptep_clear() to let page_table_check track the pte clearing operation,
    otherwise it may cause false positive in subsequent set_pte_at().

    Link: https://lkml.kernel.org/r/20230810093241.1181142-1-qi.zheng@linux.dev
    Fixes: 42b2547137f5 ("arm64/mm: enable ARCH_SUPPORTS_PAGE_TABLE_CHECK")
    Signed-off-by: Qi Zheng <zhengqi.arch@bytedance.com>
    Acked-by: Will Deacon <will@kernel.org>
    Cc: Catalin Marinas <catalin.marinas@arm.com>
    Cc: Kefeng Wang <wangkefeng.wang@huawei.com>
    Cc: Muchun Song <muchun.song@linux.dev>
    Cc: Pasha Tatashin <pasha.tatashin@soleen.com>
    Cc: Qi Zheng <zhengqi.arch@bytedance.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f428d1e8ee7209c886b00af0eadc73c2e41e42a5
Author: Jinjie Ruan <ruanjinjie@huawei.com>
Date:   Wed Aug 23 11:46:37 2023 +0800

    Bluetooth: btusb: Do not call kfree_skb() under spin_lock_irqsave()
```

```
    [ Upstream commit 2a05334d7f91ff189692089c05fc48cc1d8204de ]

    It is not allowed to call kfree_skb() from hardware interrupt
    context or with hardware interrupts being disabled.
    So replace kfree_skb() with dev_kfree_skb_irq() under
    spin_lock_irqsave(). Compile tested only.

    Fixes: baac6276c0a9 ("Bluetooth: btusb: handle mSBC audio over USB Endpoints")
    Signed-off-by: Jinjie Ruan <ruanjinjie@huawei.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e94b898463a62b72a2a8b75dea8936bf4db78e00
Author: Pauli Virtanen <pav@iki.fi>
Date:   Sat Aug 19 16:33:36 2023 +0300

    Bluetooth: hci_conn: fail SCO/ISO via hci_conn_failed if ACL gone early

    [ Upstream commit 3344d318337d9dca928fd448e966557ec5063f85 ]

    Not calling hci_(dis)connect_cfm before deleting conn referred to by a
    socket generally results to use-after-free.

    When cleaning up SCO connections when the parent ACL is deleted too
    early, use hci_conn_failed to do the connection cleanup properly.

    We also need to clean up ISO connections in a similar situation when
    connecting has started but LE Create CIS is not yet sent, so do it too
    here.

    Fixes: ca1fd42e7dbf ("Bluetooth: Fix potential double free caused by hci_conn_unlink")
    Reported-by: syzbot+cf54c1da6574b6c1b049@syzkaller.appspotmail.com
    Closes: https://lore.kernel.org/linux-bluetooth/00000000000013b93805fbbadc50@google.com/
    Signed-off-by: Pauli Virtanen <pav@iki.fi>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ba3ba53ce1f76fc372b8f918fece4f9b1e41acd4
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Aug 9 16:49:33 2023 -0700

    Bluetooth: hci_sync: Fix UAF in hci_disconnect_all_sync

    [ Upstream commit 94d9ba9f9888b748d4abd2aa1547af56ae85f772 ]

    Use-after-free can occur in hci_disconnect_all_sync if a connection is
    deleted by concurrent processing of a controller event.

    To prevent this the code now tries to iterate over the list backwards
    to ensure the links are cleanup before its parents, also it no longer
    relies on a cursor, instead it always uses the last element since
    hci_abort_conn_sync is guaranteed to call hci_conn_del.

    UAF crash log:
    ==================================================================
    BUG: KASAN: slab-use-after-free in hci_set_powered_sync
    (net/bluetooth/hci_sync.c:5424) [bluetooth]
    Read of size 8 at addr ffff888009d9c000 by task kworker/u9:0/124

    CPU: 0 PID: 124 Comm: kworker/u9:0 Tainted: G        W
    6.5.0-rc1+ #10
    Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS
    1.16.2-1.fc38 04/01/2014
    Workqueue: hci0 hci_cmd_sync_work [bluetooth]
```

```
      Call Trace:
       <TASK>
       dump_stack_lvl+0x5b/0x90
       print_report+0xcf/0x670
       ? __virt_addr_valid+0xdd/0x160
       ? hci_set_powered_sync+0x2c9/0x4a0 [bluetooth]
       kasan_report+0xa6/0xe0
       ? hci_set_powered_sync+0x2c9/0x4a0 [bluetooth]
       ? __pfx_set_powered_sync+0x10/0x10 [bluetooth]
       hci_set_powered_sync+0x2c9/0x4a0 [bluetooth]
       ? __pfx_hci_set_powered_sync+0x10/0x10 [bluetooth]
       ? __pfx_lock_release+0x10/0x10
       ? __pfx_set_powered_sync+0x10/0x10 [bluetooth]
       hci_cmd_sync_work+0x137/0x220 [bluetooth]
       process_one_work+0x526/0x9d0
       ? __pfx_process_one_work+0x10/0x10
       ? __pfx_do_raw_spin_lock+0x10/0x10
       ? mark_held_locks+0x1a/0x90
       worker_thread+0x92/0x630
       ? __pfx_worker_thread+0x10/0x10
       kthread+0x196/0x1e0
       ? __pfx_kthread+0x10/0x10
       ret_from_fork+0x2c/0x50
       </TASK>

      Allocated by task 1782:
       kasan_save_stack+0x33/0x60
       kasan_set_track+0x25/0x30
       __kasan_kmalloc+0x8f/0xa0
       hci_conn_add+0xa5/0xa80 [bluetooth]
       hci_bind_cis+0x881/0x9b0 [bluetooth]
       iso_connect_cis+0x121/0x520 [bluetooth]
       iso_sock_connect+0x3f6/0x790 [bluetooth]
       __sys_connect+0x109/0x130
       __x64_sys_connect+0x40/0x50
       do_syscall_64+0x60/0x90
       entry_SYSCALL_64_after_hwframe+0x6e/0xd8

      Freed by task 695:
       kasan_save_stack+0x33/0x60
       kasan_set_track+0x25/0x30
       kasan_save_free_info+0x2b/0x50
       __kasan_slab_free+0x10a/0x180
       __kmem_cache_free+0x14d/0x2e0
       device_release+0x5d/0xf0
       kobject_put+0xdf/0x270
       hci_disconn_complete_evt+0x274/0x3a0 [bluetooth]
       hci_event_packet+0x579/0x7e0 [bluetooth]
       hci_rx_work+0x287/0xaa0 [bluetooth]
       process_one_work+0x526/0x9d0
       worker_thread+0x92/0x630
       kthread+0x196/0x1e0
       ret_from_fork+0x2c/0x50
      ================================================================

      Fixes: 182ee45da083 ("Bluetooth: hci_sync: Rework hci_suspend_notifier")
      Signed-off-by: Pauli Virtanen <pav@iki.fi>
      Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 80265dd1d944c3f33e52375b5dbe654980bd2688
  Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
  Date:   Thu Aug 3 11:04:51 2023 -0700

      Bluetooth: hci_sync: Fix UAF on hci_abort_conn_sync
```

```
    [ Upstream commit 5af1f84ed13a416297ab9ced7537f4d5ae7f329a ]

    Connections may be cleanup while waiting for the commands to complete so
    this attempts to check if the connection handle remains valid in case of
    errors that would lead to call hci_conn_failed:

    BUG: KASAN: slab-use-after-free in hci_conn_failed+0x1f/0x160
    Read of size 8 at addr ffff888001376958 by task kworker/u3:0/52

    CPU: 0 PID: 52 Comm: kworker/u3:0 Not tainted
    6.5.0-rc1-00527-g2dfe76d58d3a #5615
    Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS
    1.16.2-1.fc38 04/01/2014
    Workqueue: hci0 hci_cmd_sync_work
    Call Trace:
     <TASK>
     dump_stack_lvl+0x1d/0x70
     print_report+0xce/0x620
     ? __virt_addr_valid+0xd4/0x150
     ? hci_conn_failed+0x1f/0x160
     kasan_report+0xd1/0x100
     ? hci_conn_failed+0x1f/0x160
     hci_conn_failed+0x1f/0x160
     hci_abort_conn_sync+0x237/0x360

    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: 94d9ba9f9888 ("Bluetooth: hci_sync: Fix UAF in hci_disconnect_all_sync")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 993fffbcc6164a9b9b6446f21f3caa649e3c7346
Author: Iulia Tanasescu <iulia.tanasescu@nxp.com>
Date:   Mon Jul 3 10:02:38 2023 +0300

    Bluetooth: ISO: Notify user space about failed bis connections

    [ Upstream commit f777d88278170410b06a1f6633f3b9375a4ddd6b ]

    Some use cases require the user to be informed if BIG synchronization
    fails. This commit makes it so that even if the BIG sync established
    event arrives with error status, a new hconn is added for each BIS,
    and the iso layer is notified about the failed connections.

    Unsuccesful bis connections will be marked using the
    HCI_CONN_BIG_SYNC_FAILED flag. From the iso layer, the POLLERR event
    is triggered on the newly allocated bis sockets, before adding them
    to the accept list of the parent socket.

    From user space, a new fd for each failed bis connection will be
    obtained by calling accept. The user should check for the POLLERR
    event on the new socket, to determine if the connection was successful
    or not.

    The HCI_CONN_BIG_SYNC flag has been added to mark whether the BIG sync
    has been successfully established. This flag is checked at bis cleanup,
    so the HCI LE BIG Terminate Sync command is only issued if needed.

    The BT_SK_BIG_SYNC flag indicates if BIG create sync has been called
    for a listening socket, to avoid issuing the command everytime a BIGInfo
    advertising report is received.

    Signed-off-by: Iulia Tanasescu <iulia.tanasescu@nxp.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: 94d9ba9f9888 ("Bluetooth: hci_sync: Fix UAF in hci_disconnect_all_sync")
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 4ab81f16c68a602b2b69e333ae08d8748a9398de
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Mon Jun 26 17:25:06 2023 -0700

    Bluetooth: hci_conn: Consolidate code for aborting connections

    [ Upstream commit a13f316e90fdb1fb6df6582e845aa9b3270f3581 ]

    This consolidates code for aborting connections using
    hci_cmd_sync_queue so it is synchronized with other threads, but
    because of the fact that some commands may block the cmd_sync_queue
    while waiting specific events this attempt to cancel those requests by
    using hci_cmd_sync_cancel.

    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: 94d9ba9f9888 ("Bluetooth: hci_sync: Fix UAF in hci_disconnect_all_sync")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c2509f7c37355e1f0bd5b7087815b845fd383723
Author: Dmitry Antipov <dmantipov@yandex.ru>
Date:   Mon Aug 14 12:49:57 2023 +0300

    wifi: mwifiex: avoid possible NULL skb pointer dereference

    [ Upstream commit 35a7a1ce7c7d61664ee54f5239a1f120ab95a87e ]

    In 'mwifiex_handle_uap_rx_forward()', always check the value
    returned by 'skb_copy()' to avoid potential NULL pointer
    dereference in 'mwifiex_uap_queue_bridged_pkt()', and drop
    original skb in case of copying failure.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Fixes: 838e4f449297 ("mwifiex: improve uAP RX handling")
    Acked-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Dmitry Antipov <dmantipov@yandex.ru>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230814095041.16416-1-dmantipov@yandex.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bdbcd4e50f2780c74790055a1a84f206d17264ac
Author: Kumar Kartikeya Dwivedi <memxor@gmail.com>
Date:   Tue Aug 22 23:21:39 2023 +0530

    bpf: Fix check_func_arg_reg_off bug for graph root/node

    [ Upstream commit 6785b2edf48c6b1c3ea61fe3b0d2e02b8fbf90c0 ]

    The commit being fixed introduced a hunk into check_func_arg_reg_off
    that bypasses reg->off == 0 enforcement when offset points to a graph
    node or root. This might possibly be done for treating bpf_rbtree_remove
    and others as KF_RELEASE and then later check correct reg->off in helper
    argument checks.

    But this is not the case, those helpers are already not KF_RELEASE and
    permit non-zero reg->off and verify it later to match the subobject in
    BTF type.

    However, this logic leads to bpf_obj_drop permitting free of register
    arguments with non-zero offset when they point to a graph root or node
    within them, which is not ok.

    For instance:
```

```
    struct foo {
            int i;
            int j;
            struct bpf_rb_node node;
    };

    struct foo *f = bpf_obj_new(typeof(*f));
    if (!f) ...
    bpf_obj_drop(f); // OK
    bpf_obj_drop(&f->i); // still ok from verifier PoV
    bpf_obj_drop(&f->node); // Not OK, but permitted right now

    Fix this by dropping the whole part of code altogether.

    Fixes: 6a3cd3318ff6 ("bpf: Migrate release_on_unlock logic to non-owning ref semantics")
    Signed-off-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230822175140.1317749-2-memxor@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit df07a3d87f6bc8c7e1fb8ed81dc201650116d6b5
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Fri Jun 23 17:24:00 2023 +0200

    mac80211: make ieee80211_tx_info padding explicit

    [ Upstream commit a7a2ef0c4b3efbd7d6f3fabd87dbbc0b3f2de5af ]

    While looking at a bug, I got rather confused by the layout of the
    'status' field in ieee80211_tx_info. Apparently, the intention is that
    status_driver_data[] is used for driver specific data, and fills up the
    size of the union to 40 bytes, just like the other ones.

    This is indeed what actually happens, but only because of the
    combination of two mistakes:

     - "void *status_driver_data[18 / sizeof(void *)];" is intended
       to be 18 bytes long but is actually two bytes shorter because of
       rounding-down in the division, to a multiple of the pointer
       size (4 bytes or 8 bytes).

     - The other fields combined are intended to be 22 bytes long, but
       are actually 24 bytes because of padding in front of the
       unaligned tx_time member, and in front of the pointer array.

    The two mistakes cancel out. so the size ends up fine, but it seems
    more helpful to make this explicit, by having a multiple of 8 bytes
    in the size calculation and explicitly describing the padding.

    Fixes: ea5907db2a9cc ("mac80211: fix struct ieee80211_tx_info size")
    Fixes: 02219b3abca59 ("mac80211: add WMM admission control support")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Reviewed-by: Kees Cook <keescook@chromium.org>
    Link: https://lore.kernel.org/r/20230623152443.2296825-2-arnd@kernel.org
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 97d36ece9860afcfe262b1568b00b6345daa7951
Author: Lin Ma <linma@zju.edu.cn>
Date:   Wed Aug 9 11:31:51 2023 +0800

    wifi: nl80211/cfg80211: add forgotten nla_policy for BSS color attribute

    [ Upstream commit 218d690c49b7e9c94ad0d317adbdd4af846ea0dc ]
```

```
    The previous commit dd3e4fc75b4a ("nl80211/cfg80211: add BSS color to
    NDP ranging parameters") adds a parameter for NDP ranging by introducing
    a new attribute type named NL80211_PMSR_FTM_REQ_ATTR_BSS_COLOR.

    However, the author forgot to also describe the nla_policy at
    nl80211_pmsr_ftm_req_attr_policy (net/wireless/nl80211.c). Just
    complement it to avoid malformed attribute that causes out-of-attribute
    access.

    Fixes: dd3e4fc75b4a ("nl80211/cfg80211: add BSS color to NDP ranging parameters")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Link: https://lore.kernel.org/r/20230809033151.768910-1-linma@zju.edu.cn
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 33608c758853c0bf79d6fb2ead7ef39e20722e71
Author: Johannes Berg <johannes.berg@intel.com>
Date:   Wed Aug 16 12:13:36 2023 +0200

    wifi: mac80211: fix puncturing bitmap handling in CSA

    [ Upstream commit 927521170c4a18c620f97865f7bad48f17c48967 ]

    Code inspection reveals that we switch the puncturing bitmap
    before the real channel switch, since that happens only in
    the second round of the worker after the channel context is
    switched by ieee80211_link_use_reserved_context().

    Fixes: 2cc25e4b2a04 ("wifi: mac80211: configure puncturing bitmap")
    Signed-off-by: Johannes Berg <johannes.berg@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4782968e0d631b0d8944dcfd4bf8fb49be087101
Author: Yonghong Song <yonghong.song@linux.dev>
Date:   Mon Aug 21 22:00:53 2023 -0700

    bpf: Fix a bpf_kptr_xchg() issue with local kptr

    [ Upstream commit ab6c637ad0276e42f8acabcbc64932a6d346dab3 ]

    When reviewing local percpu kptr support, Alexei discovered a bug
    wherea bpf_kptr_xchg() may succeed even if the map value kptr type and
    locally allocated obj type do not match ([1]). Missed struct btf_id
    comparison is the reason for the bug. This patch added such struct btf_id
    comparison and will flag verification failure if types do not match.

      [1] https://lore.kernel.org/bpf/20230819002907.io3iphmnuk43xblu@macbook-pro-
8.dhcp.thefacebook.com/#t

    Reported-by: Alexei Starovoitov <ast@kernel.org>
    Fixes: 738c96d5e2e3 ("bpf: Allow local kptrs to be exchanged via bpf_kptr_xchg")
    Signed-off-by: Yonghong Song <yonghong.song@linux.dev>
    Acked-by: Kumar Kartikeya Dwivedi <memxor@gmail.com>
    Link: https://lore.kernel.org/r/20230822050053.2886960-1-yonghong.song@linux.dev
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa61ffb021d8157dad15b6d9a5762c9f3f1741e3
Author: Fedor Pchelkin <pchelkin@ispras.ru>
Date:   Tue Apr 25 22:26:07 2023 +0300

    wifi: ath9k: protect WMI command response buffer replacement with a lock

    [ Upstream commit 454994cfa9e4c18b6df9f78b60db8eadc20a6c25 ]
```

```
    If ath9k_wmi_cmd() has exited with a timeout, it is possible that during
    next ath9k_wmi_cmd() call the wmi_rsp callback for previous wmi command
    writes to new wmi->cmd_rsp_buf and makes a completion. This results in an
    invalid ath9k_wmi_cmd() return value.

    Move the replacement of WMI command response buffer and length under
    wmi_lock. Note that last_seq_id value is updated there, too.

    Thus, the buffer cannot be written to by a belated wmi_rsp callback
    because that path is properly rejected by the last_seq_id check.

    Found by Linux Verification Center (linuxtesting.org) with Syzkaller.

    Fixes: fb9987d0f748 ("ath9k_htc: Support for AR9271 chipset.")
    Signed-off-by: Fedor Pchelkin <pchelkin@ispras.ru>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230425192607.18015-2-pchelkin@ispras.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit da2c8a93505885fd3048fdcf3aff181f752edb15
Author: Fedor Pchelkin <pchelkin@ispras.ru>
Date:   Tue Apr 25 22:26:06 2023 +0300

    wifi: ath9k: fix races between ath9k_wmi_cmd and ath9k_wmi_ctrl_rx

    [ Upstream commit b674fb513e2e7a514fcde287c0f73915d393fdb6 ]

    Currently, the synchronization between ath9k_wmi_cmd() and
    ath9k_wmi_ctrl_rx() is exposed to a race condition which, although being
    rather unlikely, can lead to invalid behaviour of ath9k_wmi_cmd().

    Consider the following scenario:

    CPU0                                        CPU1

    ath9k_wmi_cmd(...)
      mutex_lock(&wmi->op_mutex)
      ath9k_wmi_cmd_issue(...)
      wait_for_completion_timeout(...)
      ---
      timeout
      ---
                                                /* the callback is being processed
                                                 * before last_seq_id became zero
                                                 */
                                                ath9k_wmi_ctrl_rx(...)
                                                  spin_lock_irqsave(...)
                                                  /* wmi->last_seq_id check here
                                                   * doesn't detect timeout yet
                                                   */
                                                  spin_unlock_irqrestore(...)
      /* last_seq_id is zeroed to
       * indicate there was a timeout
       */
      wmi->last_seq_id = 0
      mutex_unlock(&wmi->op_mutex)
      return -ETIMEDOUT

    ath9k_wmi_cmd(...)
      mutex_lock(&wmi->op_mutex)
      /* the buffer is replaced with
       * another one
       */
```

```
        wmi->cmd_rsp_buf = rsp_buf
        wmi->cmd_rsp_len = rsp_len
        ath9k_wmi_cmd_issue(...)
          spin_lock_irqsave(...)
          spin_unlock_irqrestore(...)
        wait_for_completion_timeout(...)
                                              /* the continuation of the
                                               * callback left after the first
                                               * ath9k_wmi_cmd call
                                               */
                                                ath9k_wmi_rsp_callback(...)
                                                  /* copying data designated
                                                   * to already timeouted
                                                   * WMI command into an
                                                   * inappropriate wmi_cmd_buf
                                                   */
                                                  memcpy(...)
                                                  complete(&wmi->cmd_wait)
        /* awakened by the bogus callback
         * => invalid return result
         */
        mutex_unlock(&wmi->op_mutex)
        return 0
```

    To fix this, update last_seq_id on timeout path inside ath9k_wmi_cmd()
    under the wmi_lock. Move ath9k_wmi_rsp_callback() under wmi_lock inside
    ath9k_wmi_ctrl_rx() so that the wmi->cmd_wait can be completed only for
    initially designated wmi_cmd call, otherwise the path would be rejected
    with last_seq_id check.

    Found by Linux Verification Center (linuxtesting.org) with Syzkaller.

    Fixes: fb9987d0f748 ("ath9k_htc: Support for AR9271 chipset.")
    Signed-off-by: Fedor Pchelkin <pchelkin@ispras.ru>
    Acked-by: Toke Høiland-Jørgensen <toke@toke.dk>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230425192607.18015-1-pchelkin@ispras.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8968f684717ebbdab1160001c82c757b67a36957
Author: Daniel T. Lee <danieltimlee@gmail.com>
Date: Fri Aug 18 18:01:17 2023 +0900

    samples/bpf: fix broken map lookup probe

    [Upstream commit d93a7cf6ca2cfcd7de5d06f753ce8d5e863316ac]

    In the commit 7c4cd051add3 ("bpf: Fix syscall's stackmap lookup
    potential deadlock"), a potential deadlock issue was addressed, which
    resulted in *_map_lookup_elem not triggering BPF programs.
    (prior to lookup, bpf_disable_instrumentation() is used)

    To resolve the broken map lookup probe using "htab_map_lookup_elem",
    this commit introduces an alternative approach. Instead, it utilize
    "bpf_map_copy_value" and apply a filter specifically for the hash table
    with map_type.

    Signed-off-by: Daniel T. Lee <danieltimlee@gmail.com>
    Fixes: 7c4cd051add3 ("bpf: Fix syscall's stackmap lookup potential deadlock")
    Link: https://lore.kernel.org/r/20230818090119.477441-8-danieltimlee@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 670f616d6c4a26070a013ffc5cfd85f4bc2d8f6c
Author: Daniel T. Lee <danieltimlee@gmail.com>

```
Date:   Fri Aug 18 18:01:16 2023 +0900

    samples/bpf: fix bio latency check with tracepoint

    [ Upstream commit 92632115fb57ff9e368f256913e96d6fd5abf5ab ]

    Recently, a new tracepoint for the block layer, specifically the
    block_io_start/done tracepoints, was introduced in commit 5a80bd075f3b
    ("block: introduce block_io_start/block_io_done tracepoints").

    Previously, the kprobe entry used for this purpose was quite unstable
    and inherently broke relevant probes [1]. Now that a stable tracepoint
    is available, this commit replaces the bio latency check with it.

    One of the changes made during this replacement is the key used for the
    hash table. Since 'struct request' cannot be used as a hash key, the
    approach taken follows that which was implemented in bcc/biolatency [2].
    (uses dev:sector for the key)

    [1]: https://github.com/iovisor/bcc/issues/4261
    [2]: https://github.com/iovisor/bcc/pull/4691

    Fixes: 450b7879e345 ("block: move blk_account_io_{start,done} to blk-mq.c")
    Signed-off-by: Daniel T. Lee <danieltimlee@gmail.com>
    Link: https://lore.kernel.org/r/20230818090119.477441-7-danieltimlee@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a4f0d041428607be0ac9af20d18e047a4e31f823
Author: Aleksa Sarai <cyphar@cyphar.com>
Date:   Mon Aug 14 18:40:57 2023 +1000

    selftests: memfd: error out test process when child test fails

    [ Upstream commit 99f34659e78b9b781a3248e0b080b4dfca4957e2 ]

    Patch series "memfd: cleanups for vm.memfd_noexec", v2.

    The most critical issue with vm.memfd_noexec=2 (the fact that passing
    MFD_EXEC would bypass it entirely[1]) has been fixed in Andrew's
    tree[2], but there are still some outstanding issues that need to be
    addressed:

     * vm.memfd_noexec=2 shouldn't reject old-style memfd_create(2) syscalls
       because it will make it far to difficult to ever migrate. Instead it
       should imply MFD_EXEC.

     * The dmesg warnings are pr_warn_once(), which on most systems means
       that they will be used up by systemd or some other boot process and
       userspace developers will never see it.

       - For the !(flags & (MFD_EXEC | MFD_NOEXEC_SEAL)) case, outputting a
         rate-limited message to the kernel log is necessary to tell
         userspace that they should add the new flags.

         Arguably the most ideal way to deal with the spam concern[3,4]
         while still prompting userspace to switch to the new flags would be
         to only log the warning once per task or something similar.
         However, adding something to task_struct for tracking this would be
         needless bloat for a single pr_warn_ratelimited().

         So just switch to pr_info_ratelimited() to avoid spamming the log
         with something that isn't a real warning. There's lots of
         info-level stuff in dmesg, it seems really unlikely that this
         should be an actual problem. Most programs are already switching to
```

```
          the new flags anyway.

        - For the vm.memfd_noexec=2 case, we need to log a warning for every
          failure because otherwise userspace will have no idea why their
          previously working program started returning -EACCES (previously
          -EINVAL) from memfd_create(2). pr_warn_once() is simply wrong here.

      * The racheting mechanism for vm.memfd_noexec makes it incredibly
        unappealing for most users to enable the sysctl because enabling it
        on &init_pid_ns means you need a system reboot to unset it. Given the
        actual security threat being protected against, CAP_SYS_ADMIN users
        being restricted in this way makes little sense.

        The argument for this ratcheting by the original author was that it
        allows you to have a hierarchical setting that cannot be unset by
        child pidnses, but this is not accurate -- changing the parent
        pidns's vm.memfd_noexec setting to be more restrictive didn't affect
        children.

        Instead, switch the vm.memfd_noexec sysctl to be properly
        hierarchical and allow CAP_SYS_ADMIN users (in the pidns's owning
        userns) to lower the setting as long as it is not lower than the
        parent's effective setting. This change also makes it so that
        changing a parent pidns's vm.memfd_noexec will affect all
        descendants, providing a properly hierarchical setting. The
        performance impact of this is incredibly minimal since the maximum
        depth of pidns is 32 and it is only checked during memfd_create(2)
        and unshare(CLONE_NEWPID).

      * The memfd selftests would not exit with a non-zero error code when
        certain tests that ran in a forked process (specifically the ones
        related to MFD_EXEC and MFD_NOEXEC_SEAL) failed.

    [1]: https://lore.kernel.org/all/ZJwcsU0vI-nzgOB_@codewreck.org/
    [2]: https://lore.kernel.org/all/20230705063315.3680666-1-jeffxu@google.com/
    [3]: https://lore.kernel.org/Y5yS8wCnuYGLHMj4@x1n/
    [4]: https://lore.kernel.org/f185bb42-b29c-977e-312e-3349eea15383@linuxfoundation.org/

    This patch (of 5):

    Before this change, a test runner using this self test would see a return
    code of 0 when the tests using a child process (namely the MFD_NOEXEC_SEAL
    and MFD_EXEC tests) failed, masking test failures.

    Link: https://lkml.kernel.org/r/20230814-memfd-vm-noexec-uapi-fixes-v2-0-
7ff9e3e10ba6@cyphar.com
    Link: https://lkml.kernel.org/r/20230814-memfd-vm-noexec-uapi-fixes-v2-1-
7ff9e3e10ba6@cyphar.com
    Fixes: 11f75a01448f ("selftests/memfd: add tests for MFD_NOEXEC_SEAL MFD_EXEC")
    Signed-off-by: Aleksa Sarai <cyphar@cyphar.com>
    Reviewed-by: Jeff Xu <jeffxu@google.com>
    Cc: "Christian Brauner (Microsoft)" <brauner@kernel.org>
    Cc: Daniel Verkamp <dverkamp@chromium.org>
    Cc: Dominique Martinet <asmadeus@codewreck.org>
    Cc: Kees Cook <keescook@chromium.org>
    Cc: Shuah Khan <shuah@kernel.org>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 63940d2b3916ecad68707c5708b5ee4233f37391
Author: Ping-Ke Shih <pkshih@realtek.com>
Date:   Thu Aug 3 19:01:50 2023 +0800

    wifi: rtw89: 8852b: rfk: fine tune IQK parameters to improve performance on 2GHz band
```

```
    [ Upstream commit b3bfc4fb1edc8136396ece2d7204c2ee5cae188d ]

    A few samples get bad performance on 2GHz band, so use proper IQK command
    code and select another group to have wider range of calibration value.

    Fixes: f2abe804e823 ("wifi: rtw89: 8852b: rfk: add IQK")
    Signed-off-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230803110150.8457-1-pkshih@realtek.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7be3bf0bb2990e40fd9d1b2da317e66d57a7a3fd
Author: Polaris Pi <pinkperfect2021@gmail.com>
Date:   Thu Aug 10 08:39:11 2023 +0000

    wifi: mwifiex: Fix missed return in oob checks failed path

    [ Upstream commit 2785851c627f2db05f9271f7f63661b5dbd95c4c ]

    Add missed return in mwifiex_uap_queue_bridged_pkt() and
    mwifiex_process_rx_packet().

    Fixes: 119585281617 ("wifi: mwifiex: Fix OOB and integer underflow when rx packets")
    Signed-off-by: Polaris Pi <pinkperfect2021@gmail.com>
    Reported-by: Dmitry Antipov <dmantipov@yandex.ru>
    Acked-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230810083911.3725248-1-pinkperfect2021@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f76e1da838377777557d78dfeb6d8c532f7118be
Author: Dmitry Antipov <dmantipov@yandex.ru>
Date:   Wed Aug 2 19:07:15 2023 +0300

    wifi: mwifiex: fix memory leak in mwifiex_histogram_read()

    [ Upstream commit 9c8fd72a5c2a031cbc680a2990107ecd958ffcdb ]

    Always free the zeroed page on return from 'mwifiex_histogram_read()'.

    Fixes: cbf6e05527a7 ("mwifiex: add rx histogram statistics support")

    Acked-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Dmitry Antipov <dmantipov@yandex.ru>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230802160726.85545-1-dmantipov@yandex.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1f80adf847e88c55df07a375d193ab2fc1c8c880
Author: Eugene Shalygin <eugene.shalygin@gmail.com>
Date:   Mon Aug 21 13:52:35 2023 +0200

    hwmon: (asus-ec-sensosrs) fix mutex path for X670E Hero

    [ Upstream commit 9c53fb0ad1acaf227718ccae16e8fb8e01c05918 ]

    A user reported that they observe race condition warning [1] and after
    looking once again into the DSDT source it was found that wrong mutex
    was used.

    [1] https://github.com/zeule/asus-ec-sensors/issues/43

    Fixes: 790dec13c012 ("hwmon: (asus-ec-sensors) add ROG Crosshair X670E Hero.")
    Signed-off-by: Eugene Shalygin <eugene.shalygin@gmail.com>
    Link: https://lore.kernel.org/r/20230821115418.25733-2-eugene.shalygin@gmail.com
```

```
    Signed-off-by: Guenter Roeck <linux@roeck-us.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 73ccc8473082a54826c9fd5262d7b98cd4cb360c
Author: Eric Dumazet <edumazet@google.com>
Date:   Sat Aug 19 04:06:46 2023 +0000

    net: annotate data-races around sk->sk_lingertime

    [ Upstream commit bc1fb82ae11753c5dec53c667a055dc37796dbd2 ]

    sk_getsockopt() runs locklessly. This means sk->sk_lingertime
    can be read while other threads are changing its value.

    Other reads also happen without socket lock being held,
    and must be annotated.

    Remove preprocessor logic using BITS_PER_LONG, compilers
    are smart enough to figure this by themselves.

    v2: fixed a clang W=1 (-Wtautological-constant-out-of-range-compare) warning
        (Jakub)

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2742e33592b490570d9dfeed16bdab2dde177a59
Author: Ruan Jinjie <ruanjinjie@huawei.com>
Date:   Fri Aug 18 13:05:04 2023 +0800

    net: lan966x: Fix return value check for vcap_get_rule()

    [ Upstream commit ab104318f63997113b0ce7ac288e51359925ed79 ]

    As Simon Horman suggests, update vcap_get_rule() to always
    return an ERR_PTR() and update the error detection conditions to
    use IS_ERR(), so use IS_ERR() to fix the return value issue.

    Fixes: 72df3489fb10 ("net: lan966x: Add ptp trap rules")
    Signed-off-by: Ruan Jinjie <ruanjinjie@huawei.com>
    Suggested-by: Simon Horman <horms@kernel.org>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed2cfd57b3c12c1c437c3b1e77f824834262311c
Author: Artem Chernyshev <artem.chernyshev@red-soft.ru>
Date:   Thu Aug 3 17:54:17 2023 +0300

    fs: ocfs2: namei: check return value of ocfs2_add_entry()

    [ Upstream commit 6b72e5f9e79360fce4f2be7fe81159fbdf4256a5 ]

    Process result of ocfs2_add_entry() in case we have an error
    value.

    Found by Linux Verification Center (linuxtesting.org) with SVACE.

    Link: https://lkml.kernel.org/r/20230803145417.177649-1-artem.chernyshev@red-soft.ru
    Fixes: ccd979bdbce9 ("[PATCH] OCFS2: The Second Oracle Cluster Filesystem")
    Signed-off-by: Artem Chernyshev <artem.chernyshev@red-soft.ru>
    Reviewed-by: Joseph Qi <joseph.qi@linux.alibaba.com>
    Cc: Artem Chernyshev <artem.chernyshev@red-soft.ru>
```

```
    Cc: Joel Becker <jlbec@evilplan.org>
    Cc: Kurt Hackel <kurt.hackel@oracle.com>
    Cc: Mark Fasheh <mark@fasheh.com>
    Cc: Junxiao Bi <junxiao.bi@oracle.com>
    Cc: Changwei Ge <gechangwei@live.cn>
    Cc: Gang He <ghe@suse.com>
    Cc: Jun Piao <piaojun@huawei.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8fe12e39e40ce3470e09aff1cb8906c01d519d2e
Author: Douglas Anderson <dianders@chromium.org>
Date:   Fri Aug 4 07:00:43 2023 -0700

    watchdog/hardlockup: avoid large stack frames in watchdog_hardlockup_check()

    [ Upstream commit 1f38c86bb29f4548b8df01b47a313518e6ed2dfe ]

    After commit 77c12fc95980 ("watchdog/hardlockup: add a "cpu" param to
    watchdog_hardlockup_check()") we started storing a `struct cpumask` on the
    stack in watchdog_hardlockup_check().  On systems with CONFIG_NR_CPUS set
    to 8192 this takes up 1K on the stack.  That triggers warnings with
    `CONFIG_FRAME_WARN` set to 1024.

    We'll use the new trigger_allbutcpu_cpu_backtrace() to avoid needing to
    use a CPU mask at all.

    Link:
https://lkml.kernel.org/r/20230804065935.v4.2.I501ab68cb926ee33a7c87e063d207abf09b9943c@chan
geid
    Fixes: 77c12fc95980 ("watchdog/hardlockup: add a "cpu" param to
watchdog_hardlockup_check()")
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/r/202307310955.pLZDhpnl-lkp@intel.com
    Acked-by: Michal Hocko <mhocko@suse.com>
    Reviewed-by: Petr Mladek <pmladek@suse.com>
    Cc: Lecopzer Chen <lecopzer.chen@mediatek.com>
    Cc: Pingfan Liu <kernelfans@gmail.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4d880fc4a22cf6d500dc4021c65a3b9261a4ba9c
Author: Douglas Anderson <dianders@chromium.org>
Date:   Fri Aug 4 07:00:42 2023 -0700

    nmi_backtrace: allow excluding an arbitrary CPU

    [ Upstream commit 8d539b84f1e3478436f978ceaf55a0b6cab497b5 ]

    The APIs that allow backtracing across CPUs have always had a way to
    exclude the current CPU.  This convenience means callers didn't need to
    find a place to allocate a CPU mask just to handle the common case.

    Let's extend the API to take a CPU ID to exclude instead of just a
    boolean.  This isn't any more complex for the API to handle and allows the
    hardlockup detector to exclude a different CPU (the one it already did a
    trace for) without needing to find space for a CPU mask.

    Arguably, this new API also encourages safer behavior.  Specifically if
    the caller wants to avoid tracing the current CPU (maybe because they
    already traced the current CPU) this makes it more obvious to the caller
    that they need to make sure that the current CPU ID can't change.

    [akpm@linux-foundation.org: fix trigger_allbutcpu_cpu_backtrace() stub]
```

```
    Link:
https://lkml.kernel.org/r/20230804065935.v4.1.Ia35521b91fc781368945161d7b28538f9996c182@chan
geid
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Acked-by: Michal Hocko <mhocko@suse.com>
    Cc: kernel test robot <lkp@intel.com>
    Cc: Lecopzer Chen <lecopzer.chen@mediatek.com>
    Cc: Petr Mladek <pmladek@suse.com>
    Cc: Pingfan Liu <kernelfans@gmail.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Stable-dep-of: 1f38c86bb29f ("watchdog/hardlockup: avoid large stack frames in
watchdog_hardlockup_check()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 103c52b20000ef80079102ab105ff433e3353ac3
Author: Kuan-Ying Lee <Kuan-Ying.Lee@mediatek.com>
Date:   Mon Jul 10 17:28:46 2023 +0800

    scripts/gdb: fix 'lx-lsmod' show the wrong size

    [ Upstream commit fb40b0537342e1acd5c2daf2ff6780c1d0d2883c ]

    'lsmod' shows total core layout size, so we need to sum up all the
    sections in core layout in gdb scripts.

    / # lsmod
    kasan_test 200704 0 - Live 0xffff80007f640000

    Before patch:
    (gdb) lx-lsmod
    Address            Module                  Size  Used by
    0xffff80007f640000 kasan_test             36864  0

    After patch:
    (gdb) lx-lsmod
    Address            Module                  Size  Used by
    0xffff80007f640000 kasan_test            200704  0

    Link: https://lkml.kernel.org/r/20230710092852.31049-1-Kuan-Ying.Lee@mediatek.com
    Fixes: b4aff7513df3 ("scripts/gdb: use mem instead of core_layout to get the module
address")
    Signed-off-by: Kuan-Ying Lee <Kuan-Ying.Lee@mediatek.com>
    Reviewed-by: Pankaj Raghav <p.raghav@samsung.com>
    Cc: AngeloGioacchino Del Regno <angelogioacchino.delregno@collabora.com>
    Cc: Chinwen Chang <chinwen.chang@mediatek.com>
    Cc: Jan Kiszka <jan.kiszka@siemens.com>
    Cc: Kieran Bingham <kbingham@kernel.org>
    Cc: Luis Chamberlain <mcgrof@kernel.org>
    Cc: Matthias Brugger <matthias.bgg@gmail.com>
    Cc: Qun-Wei Lin <qun-wei.lin@mediatek.com>
    Signed-off-by: Andrew Morton <akpm@linux-foundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b9fbfb349eacc0820f91c797d7f0a3ac7a4935b5
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 4 15:14:14 2023 -0400

    USB: core: Fix race by not overwriting udev->descriptor in hub_port_init()

    commit ff33299ec8bb80cdcc073ad9c506bd79bb2ed20b upstream.

    Syzbot reported an out-of-bounds read in sysfs.c:read_descriptors():

    BUG: KASAN: slab-out-of-bounds in read_descriptors+0x263/0x280
drivers/usb/core/sysfs.c:883
```

```
    Read of size 8 at addr ffff88801e78b8c8 by task udevd/5011

    CPU: 0 PID: 5011 Comm: udevd Not tainted 6.4.0-rc6-syzkaller-00195-g40f71e7cd3c6 #0
    Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google
05/27/2023
    Call Trace:
     <TASK>
     __dump_stack lib/dump_stack.c:88 [inline]
     dump_stack_lvl+0xd9/0x150 lib/dump_stack.c:106
     print_address_description.constprop.0+0x2c/0x3c0 mm/kasan/report.c:351
     print_report mm/kasan/report.c:462 [inline]
     kasan_report+0x11c/0x130 mm/kasan/report.c:572
     read_descriptors+0x263/0x280 drivers/usb/core/sysfs.c:883
    ...
    Allocated by task 758:
    ...
     __do_kmalloc_node mm/slab_common.c:966 [inline]
     __kmalloc+0x5e/0x190 mm/slab_common.c:979
     kmalloc include/linux/slab.h:563 [inline]
     kzalloc include/linux/slab.h:680 [inline]
     usb_get_configuration+0x1f7/0x5170 drivers/usb/core/config.c:887
     usb_enumerate_device drivers/usb/core/hub.c:2407 [inline]
     usb_new_device+0x12b0/0x19d0 drivers/usb/core/hub.c:2545

    As analyzed by Khazhy Kumykov, the cause of this bug is a race between
    read_descriptors() and hub_port_init(): The first routine uses a field
    in udev->descriptor, not expecting it to change, while the second
    overwrites it.

    Prior to commit 45bf39f8df7f ("USB: core: Don't hold device lock while
    reading the "descriptors" sysfs file") this race couldn't occur,
    because the routines were mutually exclusive thanks to the device
    locking.  Removing that locking from read_descriptors() exposed it to
    the race.

    The best way to fix the bug is to keep hub_port_init() from changing
    udev->descriptor once udev has been initialized and registered.
    Drivers expect the descriptors stored in the kernel to be immutable;
    we should not undermine this expectation.  In fact, this change should
    have been made long ago.

    So now hub_port_init() will take an additional argument, specifying a
    buffer in which to store the device descriptor it reads.  (If udev has
    not yet been initialized, the buffer pointer will be NULL and then
    hub_port_init() will store the device descriptor in udev as before.)
    This eliminates the data race responsible for the out-of-bounds read.

    The changes to hub_port_init() appear more extensive than they really
    are, because of indentation changes resulting from an attempt to avoid
    writing to other parts of the usb_device structure after it has been
    initialized.  Similar changes should be made to the code that reads
    the BOS descriptor, but that can be handled in a separate patch later
    on.  This patch is sufficient to fix the bug found by syzbot.

    Reported-and-tested-by: syzbot+18996170f8096c6174d0@syzkaller.appspotmail.com
    Closes: https://lore.kernel.org/linux-usb/000000000000c0ffe505fe86c9ca@google.com/#r
    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Cc: Khazhy Kumykov <khazhy@google.com>
    Fixes: 45bf39f8df7f ("USB: core: Don't hold device lock while reading the "descriptors"
sysfs file")
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/b958b47a-9a46-4c22-a9f9-e42e42c31251@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 91eafd344f7467d3e078d5ac5d8084f32e4e166d
```

```
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 4 15:12:21 2023 -0400

    USB: core: Change usb_get_device_descriptor() API

    commit de28e469da75359a2bb8cd8778b78aa64b1be1f4 upstream.

    The usb_get_device_descriptor() routine reads the device descriptor
    from the udev device and stores it directly in udev->descriptor.  This
    interface is error prone, because the USB subsystem expects in-memory
    copies of a device's descriptors to be immutable once the device has
    been initialized.

    The interface is changed so that the device descriptor is left in a
    kmalloc-ed buffer, not copied into the usb_device structure.  A
    pointer to the buffer is returned to the caller, who is then
    responsible for kfree-ing it.  The corresponding changes needed in the
    various callers are fairly small.

    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Link: https://lore.kernel.org/r/d0111bb6-56c1-4f90-adf2-6cfe152f6561@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit a2bd8ade86d9be6f7fbb22a2968e02964a3daf39
Author: Alan Stern <stern@rowland.harvard.edu>
Date:   Fri Aug 4 15:10:59 2023 -0400

    USB: core: Unite old scheme and new scheme descriptor reads

    commit 85d07c55621676d47d873d2749b88f783cd4d5a1 upstream.

    In preparation for reworking the usb_get_device_descriptor() routine,
    it is desirable to unite the two different code paths responsible for
    initially determining endpoint 0's maximum packet size in a newly
    discovered USB device.  Making this determination presents a
    chicken-and-egg sort of problem, in that the only way to learn the
    maxpacket value is to get it from the device descriptor retrieved from
    the device, but communicating with the device to retrieve a descriptor
    requires us to know beforehand the ep0 maxpacket size.

    In practice this problem is solved in two different ways, referred to
    in hub.c as the "old scheme" and the "new scheme".  The old scheme
    (which is the approach recommended by the USB-2 spec) involves asking
    the device to send just the first eight bytes of its device
    descriptor.  Such a transfer uses packets containing no more than
    eight bytes each, and every USB device must have an ep0 maxpacket size
    >= 8, so this should succeed.  Since the bMaxPacketSize0 field of the
    device descriptor lies within the first eight bytes, this is all we
    need.

    The new scheme is an imitation of the technique used in an early
    Windows USB implementation, giving it the happy advantage of working
    with a wide variety of devices (some of them at the time would not
    work with the old scheme, although that's probably less true now).  It
    involves making an initial guess of the ep0 maxpacket size, asking the
    device to send up to 64 bytes worth of its device descriptor (which is
    only 18 bytes long), and then resetting the device to clear any error
    condition that might have resulted from the guess being wrong.  The
    initial guess is determined by the connection speed; it should be
    correct in all cases other than full speed, for which the allowed
    values are 8, 16, 32, and 64 (in this case the initial guess is 64).

    The reason for this patch is that the old- and new-scheme parts of
    hub_port_init() use different code paths, one involving
    usb_get_device_descriptor() and one not, for their initial reads of
```

```
    the device descriptor.  Since these reads have essentially the same
    purpose and are made under essentially the same circumstances, this is
    illogical.  It makes more sense to have both of them use a common
    subroutine.

    This subroutine does basically what the new scheme's code did, because
    that approach is more general than the one used by the old scheme.  It
    only needs to know how many bytes to transfer and whether or not it is
    being called for the first iteration of a retry loop (in case of
    certain time-out errors).  There are two main differences from the
    former code:

            We initialize the bDescriptorType field of the transfer buffer
            to 0 before performing the transfer, to avoid possibly
            accessing an uninitialized value afterward.

            We read the device descriptor into a temporary buffer rather
            than storing it directly into udev->descriptor, which the old
            scheme implementation used to do.

    Since the whole point of this first read of the device descriptor is
    to determine the bMaxPacketSize0 value, that is what the new routine
    returns (or an error code).  The value is stored in a local variable
    rather than in udev->descriptor.  As a side effect, this necessitates
    moving a section of code that checks the bcdUSB field for SuperSpeed
    devices until after the full device descriptor has been retrieved.

    Signed-off-by: Alan Stern <stern@rowland.harvard.edu>
    Cc: Oliver Neukum <oneukum@suse.com>
    Link: https://lore.kernel.org/r/495cb5d4-f956-4f4a-a875-1e67e9489510@rowland.harvard.edu
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 1101867a1711c27d8bbe0e83136bec47f8c1ca2a
Author: RD Babiera <rdbabiera@google.com>
Date:   Mon Aug 14 18:05:59 2023 +0000

    usb: typec: bus: verify partner exists in typec_altmode_attention

    commit f23643306430f86e2f413ee2b986e0773e79da31 upstream.

    Some usb hubs will negotiate DisplayPort Alt mode with the device
    but will then negotiate a data role swap after entering the alt
    mode. The data role swap causes the device to unregister all alt
    modes, however the usb hub will still send Attention messages
    even after failing to reregister the Alt Mode. type_altmode_attention
    currently does not verify whether or not a device's altmode partner
    exists, which results in a NULL pointer error when dereferencing
    the typec_altmode and typec_altmode_ops belonging to the altmode
    partner.

    Verify the presence of a device's altmode partner before sending
    the Attention message to the Alt Mode driver.

    Fixes: 8a37d87d72f0 ("usb: typec: Bus type for alternate modes")
    Cc: stable@vger.kernel.org
    Signed-off-by: RD Babiera <rdbabiera@google.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Reviewed-by: Guenter Roeck <linux@roeck-us.net>
    Link: https://lore.kernel.org/r/20230814180559.923475-1-rdbabiera@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit cbcf107780aecf51aba68488044a416d95060b6d
Author: RD Babiera <rdbabiera@google.com>
Date:   Mon Jul 31 16:59:23 2023 +0000
```

```
    usb: typec: tcpm: set initial svdm version based on pd revision

    commit c97cd0b4b54eb42aed7f6c3c295a2d137f6d2416 upstream.

    When sending Discover Identity messages to a Port Partner that uses Power
    Delivery v2 and SVDM v1, we currently send PD v2 messages with SVDM v2.0,
    expecting the port partner to respond with its highest supported SVDM
    version as stated in Section 6.4.4.2.3 in the Power Delivery v3
    specification. However, sending SVDM v2 to some Power Delivery v2 port
    partners results in a NAK whereas sending SVDM v1 does not.

    NAK messages can be handled by the initiator (PD v3 section 6.4.4.2.5.1),
    and one solution could be to resend Discover Identity on a lower SVDM
    version if possible. But, Section 6.4.4.3 of PD v2 states that "A NAK
    response Should be taken as an indication not to retry that particular
    Command."

    Instead, we can set the SVDM version to the maximum one supported by the
    negotiated PD revision. When operating in PD v2, this obeys Section
    6.4.4.2.3, which states the SVDM field "Shall be set to zero to indicate
    Version 1.0." In PD v3, the SVDM field "Shall be set to 01b to indicate
    Version 2.0."

    Fixes: c34e85fa69b9 ("usb: typec: tcpm: Send DISCOVER_IDENTITY from dedicated work")
    Cc: stable@vger.kernel.org
    Signed-off-by: RD Babiera <rdbabiera@google.com>
    Reviewed-by: Heikki Krogerus <heikki.krogerus@linux.intel.com>
    Link: https://lore.kernel.org/r/20230731165926.1815338-1-rdbabiera@google.com
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 0d6b0ec0fb837a8cb5c68c85d299717f557b76f7
Author: Yan Zhai <yan@cloudflare.com>
Date:   Thu Aug 17 19:58:14 2023 -0700

    lwt: Check LWTUNNEL_XMIT_CONTINUE strictly

    [ Upstream commit a171fbec88a2c730b108c7147ac5e7b2f5a02b47 ]

    LWTUNNEL_XMIT_CONTINUE is implicitly assumed in ip(6)_finish_output2,
    such that any positive return value from a xmit hook could cause
    unexpected continue behavior, despite that related skb may have been
    freed. This could be error-prone for future xmit hook ops. One of the
    possible errors is to return statuses of dst_output directly.

    To make the code safer, redefine LWTUNNEL_XMIT_CONTINUE value to
    distinguish from dst_output statuses and check the continue
    condition explicitly.

    Fixes: 3a0af8fd61f9 ("bpf: BPF for lightweight tunnel infrastructure")
    Suggested-by: Dan Carpenter <dan.carpenter@linaro.org>
    Signed-off-by: Yan Zhai <yan@cloudflare.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link:
https://lore.kernel.org/bpf/96b939b85eda00e8df4f7c080f770970a4c5f698.1692326837.git.yan@clou
dflare.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65583f9e070db7bece20710cfa2e3daeb0b831d9
Author: Yan Zhai <yan@cloudflare.com>
Date:   Thu Aug 17 19:58:11 2023 -0700

    lwt: Fix return values of BPF xmit ops

    [ Upstream commit 29b22badb7a84b783e3a4fffca16f7768fb31205 ]
```

```
    BPF encap ops can return different types of positive values, such like
    NET_RX_DROP, NET_XMIT_CN, NETDEV_TX_BUSY, and so on, from function
    skb_do_redirect and bpf_lwt_xmit_reroute. At the xmit hook, such return
    values would be treated implicitly as LWTUNNEL_XMIT_CONTINUE in
    ip(6)_finish_output2. When this happens, skbs that have been freed would
    continue to the neighbor subsystem, causing use-after-free bug and
    kernel crashes.

    To fix the incorrect behavior, skb_do_redirect return values can be
    simply discarded, the same as tc-egress behavior. On the other hand,
    bpf_lwt_xmit_reroute returns useful errors to local senders, e.g. PMTU
    information. Thus convert its return values to avoid the conflict with
    LWTUNNEL_XMIT_CONTINUE.

    Fixes: 3a0af8fd61f9 ("bpf: BPF for lightweight tunnel infrastructure")
    Reported-by: Jordan Griege <jgriege@cloudflare.com>
    Suggested-by: Martin KaFai Lau <martin.lau@linux.dev>
    Suggested-by: Stanislav Fomichev <sdf@google.com>
    Signed-off-by: Yan Zhai <yan@cloudflare.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link:
https://lore.kernel.org/bpf/0d2b878186cfe215fec6b45769c1cd0591d3628d.1692326837.git.yan@clou
dflare.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 93ea1ee2ad712776f9c6096dc7960efecaea88bf
Author: Florian Fainelli <florian.fainelli@broadcom.com>
Date:   Thu Aug 10 12:22:08 2023 -0700

    hwrng: iproc-rng200 - Implement suspend and resume calls

    [ Upstream commit 8e03dd62e5be811efbf0cbeba47e79e793519105 ]

    Chips such as BCM7278 support system wide suspend/resume which will
    cause the HWRNG block to lose its state and reset to its power on reset
    register values. We need to cleanup and re-initialize the HWRNG for it
    to be functional coming out of a system suspend cycle.

    Fixes: c3577f6100ca ("hwrng: iproc-rng200 - Add support for BCM7278")
    Signed-off-by: Florian Fainelli <florian.fainelli@broadcom.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 814e0d5403c8c342a6ee8b4482394d2d07a07eeb
Author: Gaurav Jain <gaurav.jain@nxp.com>
Date:   Tue Aug 8 12:55:25 2023 +0200

    crypto: caam - fix unchecked return value error

    [ Upstream commit e30685204711a6be40dec2622606950ccd37dafe ]

    error:
    Unchecked return value (CHECKED_RETURN)
    check_return: Calling sg_miter_next without checking return value

    fix:
    added check if(!sg_miter_next)

    Fixes: 8a2a0dd35f2e ("crypto: caam - strip input zeros from RSA input buffer")
    Signed-off-by: Gaurav Jain <gaurav.jain@nxp.com>
    Signed-off-by: Meenakshi Aggarwal <meenakshi.aggarwal@nxp.com>
    Reviewed-by: Gaurav Jain <gaurav.jain@nxp.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit a07c6f111f0cd4c6f308c46b63f3ed5f84126086
Author: Przemek Kitszel <przemyslaw.kitszel@intel.com>
Date:   Tue Aug 8 17:54:15 2023 -0400

    ice: ice_aq_check_events: fix off-by-one check when filling buffer

    [ Upstream commit e1e8a142c43336e3d25bfa1cb3a4ae7d00875c48 ]

    Allow task's event buffer to be filled also in the case that it's size
    is exactly the size of the message.

    Fixes: d69ea414c9b4 ("ice: implement device flash update via devlink")
    Reviewed-by: Jacob Keller <jacob.e.keller@intel.com>
    Signed-off-by: Przemek Kitszel <przemyslaw.kitszel@intel.com>
    Reviewed-by: Simon Horman <horms@kernel.org>
    Tested-by: Pucha Himasekhar Reddy <himasekharx.reddy.pucha@intel.com> (A Contingent
worker at Intel)
    Signed-off-by: Tony Nguyen <anthony.l.nguyen@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1a28a27c836f0485e3e5625778f6bfd0da71fd2e
Author: Abel Wu <wuyun.abel@bytedance.com>
Date:   Mon Aug 14 15:09:11 2023 +0800

    net-memcg: Fix scope of sockmem pressure indicators

    [ Upstream commit ac8a52962164a50e693fa021d3564d7745b83a7f ]

    Now there are two indicators of socket memory pressure sit inside
    struct mem_cgroup, socket_pressure and tcpmem_pressure, indicating
    memory reclaim pressure in memcg->memory and ->tcpmem respectively.

    When in legacy mode (cgroupv1), the socket memory is charged into
    ->tcpmem which is independent of ->memory, so socket_pressure has
    nothing to do with socket's pressure at all. Things could be worse
    by taking socket_pressure into consideration in legacy mode, as a
    pressure in ->memory can lead to premature reclamation/throttling
    in socket.

    While for the default mode (cgroupv2), the socket memory is charged
    into ->memory, and ->tcpmem/->tcpmem_pressure are simply not used.

    So {socket,tcpmem}_pressure are only used in default/legacy mode
    respectively for indicating socket memory pressure. This patch fixes
    the pieces of code that make mixed use of both.

    Fixes: 8e8ae645249b ("mm: memcontrol: hook up vmpressure to socket pressure")
    Signed-off-by: Abel Wu <wuyun.abel@bytedance.com>
    Acked-by: Shakeel Butt <shakeelb@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3354d1898c8e510b68d515ad8553ea836e9d185b
Author: Jijie Shao <shaojijie@huawei.com>
Date:   Tue Aug 15 14:06:41 2023 +0800

    net: hns3: fix wrong rpu tln reg issue

    [ Upstream commit 36122201eeaefd78547def9681aa5d83b5a00b6a ]

    In the original RPU query command, the status register values of
    multiple RPU tunnels are accumulated by default, which is unreasonable.
    This patch Fix it by querying the specified tunnel ID.
    The tunnel number of the device can be obtained from firmware
    during initialization.
```

```
    Fixes: ddb54554fa51 ("net: hns3: add DFX registers information for ethtool -d")
    Signed-off-by: Jijie Shao <shaojijie@huawei.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d721bd9424124d8377db6da2995fb7cbe4898b7f
Author: Jijie Shao <shaojijie@huawei.com>
Date:   Tue Aug 15 14:06:39 2023 +0800

    net: hns3: Support tlv in regs data for HNS3 PF driver

    [ Upstream commit d8634b7c3f62d265fc2ecf29286aa9c5b78f969f ]

    The dump register function is being refactored.
    The second step in refactoring is to support tlv info in regs data for
    HNS3 PF driver.

    Currently, if we use "ethtool -d" to dump regs value,
    the output is as follows:
      offset1: 00 01 02 03 04 05 ...
      offset2 : 10 11 12 13 14 15 ...
      ......

    We can't get the value of a register directly.

    This patch deletes the original separator information and
    add tag_len_value information in regs data.
    ethtool can parse register data in key-value format by -d command.

    a patch will be added to the ethtool to parse regs data
    in the following format:
      reg1 : value2
      reg2 : value2
      ......

    Signed-off-by: Jijie Shao <shaojijie@huawei.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Stable-dep-of: 36122201eeae ("net: hns3: fix wrong rpu tln reg issue")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2b1fff96a297034f03466cfecda9824adafe16ed
Author: Jijie Shao <shaojijie@huawei.com>
Date:   Tue Aug 15 14:06:38 2023 +0800

    net: hns3: move dump regs function to a separate file

    [ Upstream commit 939ccd107ffcade20c9c7055a2e7ae0fd724fb72 ]

    The dump register function is being refactored.
    The first step in refactoring is put the dump regs function
    into a separate file.

    Signed-off-by: Jijie Shao <shaojijie@huawei.com>
    Reviewed-by: Leon Romanovsky <leonro@nvidia.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Stable-dep-of: 36122201eeae ("net: hns3: fix wrong rpu tln reg issue")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 003e72ecc2a784aef74eb833c0ef7d1876137407
Author: Yipeng Zou <zouyipeng@huawei.com>
Date:   Mon Aug 14 11:07:27 2023 +0800
```

```
    selftests/bpf: Clean up fmod_ret in bench_rename test script

    [ Upstream commit 83a89c4b6ae93481d3f618aba6a29d89208d26ed ]

    Running the bench_rename test script, the following error occurs:

      # ./benchs/run_bench_rename.sh
      base      :    0.819 ± 0.012M/s
      kprobe    :    0.538 ± 0.009M/s
      kretprobe :    0.503 ± 0.004M/s
      rawtp     :    0.779 ± 0.020M/s
      fentry    :    0.726 ± 0.007M/s
      fexit     :    0.691 ± 0.007M/s
      benchmark 'rename-fmodret' not found

    The bench_rename_fmodret has been removed in commit b000def2e052
    ("selftests: Remove fmod_ret from test_overhead"), thus remove it
    from the runners in the test script.

    Fixes: b000def2e052 ("selftests: Remove fmod_ret from test_overhead")
    Signed-off-by: Yipeng Zou <zouyipeng@huawei.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230814030727.3010390-1-zouyipeng@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 366d66a3c9323ccb6cab21f9eb38f74e2d970df2
Author: Yipeng Zou <zouyipeng@huawei.com>
Date:   Mon Aug 14 11:14:34 2023 +0800

    selftests/bpf: Fix repeat option when kfunc_call verification fails

    [ Upstream commit 811915db674f8daf19bb4fcb67da9017235ce26d ]

    There is no way where topts.repeat can be set to 1 when tc_test fails.
    Fix the typo where the break statement slipped by one line.

    Fixes: fb66223a244f ("selftests/bpf: add test for accessing ctx from syscall program
type")
    Signed-off-by: Yipeng Zou <zouyipeng@huawei.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Reviewed-by: Li Zetao <lizetao1@huawei.com>
    Link: https://lore.kernel.org/bpf/20230814031434.3077944-1-zouyipeng@huawei.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 4cdcdadc82ef4d87e41a27af498c1395447ac966
Author: Marco Vedovati <marco.vedovati@crowdstrike.com>
Date:   Thu Aug 10 14:43:53 2023 -0700

    libbpf: Set close-on-exec flag on gzopen

    [ Upstream commit 8e50750f122e59ea4cab4b4f696ef22b391bedc9 ]

    Enable the close-on-exec flag when using gzopen. This is especially important
    for multithreaded programs making use of libbpf, where a fork + exec could
    race with libbpf library calls, potentially resulting in a file descriptor
    leaked to the new process. This got missed in 59842c5451fe ("libbpf: Ensure
    libbpf always opens files with O_CLOEXEC").

    Fixes: 59842c5451fe ("libbpf: Ensure libbpf always opens files with O_CLOEXEC")
    Signed-off-by: Marco Vedovati <marco.vedovati@crowdstrike.com>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230810214350.106301-1-martin.kelly@crowdstrike.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 15b8fc7b560276fa15e6280642f4d6e252d11b00
```

```
Author: Vladimir Oltean <vladimir.oltean@nxp.com>
Date:   Fri Aug 11 14:53:52 2023 +0300

    net: pcs: lynx: fix lynx_pcs_link_up_sgmii() not doing anything in fixed-link mode

    [ Upstream commit 2f4503f94c5d81d1589842bfb457be466c8c670b ]

    lynx_pcs_link_up_sgmii() is supposed to update the PCS speed and duplex
    for the non-inband operating modes, and prior to the blamed commit, it
    did just that, but a mistake sneaked into the conversion and reversed
    the condition.

    It is easy for this to go undetected on platforms that also initialize
    the PCS in the bootloader, because Linux doesn't reset it (although
    maybe it should). The nature of the bug is that phylink will not touch
    the IF_MODE_HALF_DUPLEX | IF_MODE_SPEED_MSK fields when it should, and
    it will apparently keep working if the previous values set by the
    bootloader were correct.

    Fixes: c689a6528c22 ("net: pcs: lynx: update PCS driver to use neg_mode")
    Signed-off-by: Vladimir Oltean <vladimir.oltean@nxp.com>
    Reviewed-by: Russell King (Oracle) <rmk+kernel@armlinux.org.uk>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ea8b23d715051c595ea5711a92e868b47032ba86
Author: Menglong Dong <imagedong@tencent.com>
Date:   Fri Aug 11 10:55:29 2023 +0800

    net: tcp: fix unexcepted socket die when snd_wnd is 0

    [ Upstream commit e89688e3e97868451a5d05b38a9d2633d6785cd4 ]

    In tcp_retransmit_timer(), a window shrunk connection will be regarded
    as timeout if 'tcp_jiffies32 - tp->rcv_tstamp > TCP_RTO_MAX'. This is not
    right all the time.

    The retransmits will become zero-window probes in tcp_retransmit_timer()
    if the 'snd_wnd==0'. Therefore, the icsk->icsk_rto will come up to
    TCP_RTO_MAX sooner or later.

    However, the timer can be delayed and be triggered after 122877ms, not
    TCP_RTO_MAX, as I tested.

    Therefore, 'tcp_jiffies32 - tp->rcv_tstamp > TCP_RTO_MAX' is always true
    once the RTO come up to TCP_RTO_MAX, and the socket will die.

    Fix this by replacing the 'tcp_jiffies32' with '(u32)icsk->icsk_timeout',
    which is exact the timestamp of the timeout.

    However, "tp->rcv_tstamp" can restart from idle, then tp->rcv_tstamp
    could already be a long time (minutes or hours) in the past even on the
    first RTO. So we double check the timeout with the duration of the
    retransmission.

    Meanwhile, making "2 * TCP_RTO_MAX" as the timeout to avoid the socket
    dying too soon.

    Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2")
    Link: https://lore.kernel.org/netdev/CADxym3YyMiO+zMD4zj03YPM3FBi-
1LHi6gSD2XT8pyAMM096pg@mail.gmail.com/
    Signed-off-by: Menglong Dong <imagedong@tencent.com>
    Reviewed-by: Eric Dumazet <edumazet@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 98d1d2e0f95c58ef12edbf6a972041a4ddb94ebf
Author: Pauli Virtanen <pav@iki.fi>
Date:   Sat Aug 5 19:08:41 2023 +0300

    Bluetooth: hci_event: drop only unbound CIS if Set CIG Parameters fails

    [ Upstream commit 66dee21524d9ac6461ec3052652b7bc0603ee0c5 ]

    When user tries to connect a new CIS when its CIG is not configurable,
    that connection shall fail, but pre-existing connections shall not be
    affected.  However, currently hci_cc_le_set_cig_params deletes all CIS
    of the CIG on error so it doesn't work, even though controller shall not
    change CIG/CIS configuration if the command fails.

    Fix by failing on command error only the connections that are not yet
    bound, so that we keep the previous CIS configuration like the
    controller does.

    Fixes: 26afbd826ee3 ("Bluetooth: Add initial implementation of CIS connections")
    Signed-off-by: Pauli Virtanen <pav@iki.fi>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 090f3129538de2801bb582459546eeb4a8705310
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Wed Jun 28 12:15:53 2023 -0700

    Bluetooth: hci_conn: Always allocate unique handles

    [ Upstream commit 9f78191cc9f1b34c2e2afd7b554a83bf034092dd ]

    This attempts to always allocate a unique handle for connections so they
    can be properly aborted by the likes of hci_abort_conn, so this uses the
    invalid range as a pool of unset handles that way if userspace is trying
    to create multiple connections at once each will be given a unique
    handle which will be considered unset.

    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: 66dee21524d9 ("Bluetooth: hci_event: drop only unbound CIS if Set CIG
Parameters fails")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8d66f7ced51cb924bc90278d6a0a26a52877271a
Author: Manish Mandlik <mmandlik@google.com>
Date:   Fri Aug 4 11:14:45 2023 -0700

    Bluetooth: hci_sync: Avoid use-after-free in dbg for hci_add_adv_monitor()

    [ Upstream commit a2bcd2b63271a93a695fabbfbf459c603d956d48 ]

    KSAN reports use-after-free in hci_add_adv_monitor().

    While adding an adv monitor,
        hci_add_adv_monitor() calls ->
        msft_add_monitor_pattern() calls ->
        msft_add_monitor_sync() calls ->
        msft_le_monitor_advertisement_cb() calls in an error case ->
        hci_free_adv_monitor() which frees the *moniter.

    This is referenced by bt_dev_dbg() in hci_add_adv_monitor().

    Fix the bt_dev_dbg() by using handle instead of monitor->handle.

    Fixes: b747a83690c8 ("Bluetooth: hci_sync: Refactor add Adv Monitor")
```

```
    Signed-off-by: Manish Mandlik <mmandlik@google.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e153cfb17a23b3864306b2c6a4e60f0a9a63d367
Author: Douglas Anderson <dianders@chromium.org>
Date:   Fri Jun 30 15:33:15 2023 -0700

    Bluetooth: hci_sync: Don't double print name in add/remove adv_monitor

    [ Upstream commit 6f55eea116ba3646fb5fbb31de703f8cf79d8214 ]

    The hci_add_adv_monitor() hci_remove_adv_monitor() functions call
    bt_dev_dbg() to print some debug statements. The bt_dev_dbg() macro
    automatically adds in the device's name. That means that we shouldn't
    include the name in the bt_dev_dbg() calls.

    Suggested-by: Luiz Augusto von Dentz <luiz.dentz@gmail.com>
    Signed-off-by: Douglas Anderson <dianders@chromium.org>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Stable-dep-of: a2bcd2b63271 ("Bluetooth: hci_sync: Avoid use-after-free in dbg for
hci_add_adv_monitor()")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 35cc42f04bc49f0656f6840cb7451b3df6049649
Author: Min Li <lm0963hack@gmail.com>
Date:   Mon Aug 7 19:07:41 2023 +0800

    Bluetooth: Fix potential use-after-free when clear keys

    [ Upstream commit 3673952cf0c6cf81b06c66a0b788abeeb02ff3ae ]

    Similar to commit c5d2b6fa26b5 ("Bluetooth: Fix use-after-free in
    hci_remove_ltk/hci_remove_irk"). We can not access k after kfree_rcu()
    call.

    Fixes: d7d41682efc2 ("Bluetooth: Fix Suspicious RCU usage warnings")
    Signed-off-by: Min Li <lm0963hack@gmail.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d8570c4c3f2a3e51b3c8b5e6ec898364c5c03062
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Fri Aug 4 14:54:09 2023 -0700

    Bluetooth: hci_conn: Fix hci_le_set_cig_params

    [ Upstream commit a091289218202bc09d9b9caa8afcde1018584aec ]

    When running with concurrent task only one CIS was being assigned so
    this attempts to rework the way the PDU is constructed so it is handled
    later at the callback instead of in place.

    Fixes: 26afbd826ee3 ("Bluetooth: Add initial implementation of CIS connections")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a4d15c99dcb68257cb3dde261c055f64c80163d9
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Thu Aug 3 16:41:34 2023 -0700

    Bluetooth: hci_conn: Fix not allowing valid CIS ID

    [ Upstream commit f2f84a70f9d0c9a3263194ca9d82e7bc6027d356 ]
```

```
    Only the number of CIS shall be limited to 0x1f, the CIS ID in the
    other hand is up to 0xef.

    Fixes: 26afbd826ee3 ("Bluetooth: Add initial implementation of CIS connections")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa4ef8f8d7609c3ceba61f4c94d806b5ddac4133
Author: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
Date:   Thu Aug 3 14:41:46 2023 -0700

    Bluetooth: ISO: Fix not checking for valid CIG/CIS IDs

    [ Upstream commit b7f923b1ef6a2e76013089d30c9552257056360a ]

    Valid range of CIG/CIS are 0x00 to 0xEF, so this checks they are
    properly checked before attempting to use HCI_OP_LE_SET_CIG_PARAMS.

    Fixes: ccf74f2390d6 ("Bluetooth: Add BTPROTO_ISO socket type")
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c976a72469396886aeccad37899904ae8b2b8572
Author: Yuanjun Gong <ruc_gongyuanjun@163.com>
Date:   Wed Jul 26 21:30:00 2023 +0800

    Bluetooth: nokia: fix value check in nokia_bluetooth_serdev_probe()

    [ Upstream commit e8b5aed31355072faac8092ead4938ddec3111fd ]

    in nokia_bluetooth_serdev_probe(), check the return value of
    clk_prepare_enable() and return the error code if
    clk_prepare_enable() returns an unexpected value.

    Fixes: 7bb318680e86 ("Bluetooth: add nokia driver")
    Signed-off-by: Yuanjun Gong <ruc_gongyuanjun@163.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 722b0fb0f622b6348dd38fe9b8b08690f1a55910
Author: Pauli Virtanen <pav@iki.fi>
Date:   Thu Jun 1 09:34:46 2023 +0300

    Bluetooth: ISO: do not emit new LE Create CIS if previous is pending

    [ Upstream commit 7f74563e6140e42b4ffae62adbef7a65967a3f98 ]

    LE Create CIS command shall not be sent before all CIS Established
    events from its previous invocation have been processed. Currently it is
    sent via hci_sync but that only waits for the first event, but there can
    be multiple.

    Make it wait for all events, and simplify the CIS creation as follows:

    Add new flag HCI_CONN_CREATE_CIS, which is set if Create CIS has been
    sent for the connection but it is not yet completed.

    Make BT_CONNECT state to mean the connection wants Create CIS.

    On events after which new Create CIS may need to be sent, send it if
    possible and some connections need it. These events are:
    hci_connect_cis, iso_connect_cfm, hci_cs_le_create_cis,
    hci_le_cis_estabilished_evt.

    The Create CIS status/completion events shall queue new Create CIS only
```

```
    if at least one of the connections transitions away from BT_CONNECT, so
    that we don't loop if controller is sending bogus events.

    This fixes sending multiple CIS Create for the same CIS in the
    "ISO AC 6(i) - Success" BlueZ test case:

    < HCI Command: LE Create Co.. (0x08|0x0064) plen 9  #129 [hci0]
            Number of CIS: 2
            CIS Handle: 257
            ACL Handle: 42
            CIS Handle: 258
            ACL Handle: 42
    > HCI Event: Command Status (0x0f) plen 4           #130 [hci0]
          LE Create Connected Isochronous Stream (0x08|0x0064) ncmd 1
            Status: Success (0x00)
    > HCI Event: LE Meta Event (0x3e) plen 29           #131 [hci0]
          LE Connected Isochronous Stream Established (0x19)
            Status: Success (0x00)
            Connection Handle: 257
            ...
    < HCI Command: LE Setup Is.. (0x08|0x006e) plen 13  #132 [hci0]
            ...
    > HCI Event: Command Complete (0x0e) plen 6         #133 [hci0]
          LE Setup Isochronous Data Path (0x08|0x006e) ncmd 1
            ...
    < HCI Command: LE Create Co.. (0x08|0x0064) plen 5  #134 [hci0]
            Number of CIS: 1
            CIS Handle: 258
            ACL Handle: 42
    > HCI Event: Command Status (0x0f) plen 4           #135 [hci0]
          LE Create Connected Isochronous Stream (0x08|0x0064) ncmd 1
            Status: ACL Connection Already Exists (0x0b)
    > HCI Event: LE Meta Event (0x3e) plen 29           #136 [hci0]
          LE Connected Isochronous Stream Established (0x19)
            Status: Success (0x00)
            Connection Handle: 258
            ...

    Fixes: c09b80be6ffc ("Bluetooth: hci_conn: Fix not waiting for
HCI_EVT_LE_CIS_ESTABLISHED")
    Signed-off-by: Pauli Virtanen <pav@iki.fi>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2ccde10127447c1a5caad8469fede945bdb62fdf
Author: Iulia Tanasescu <iulia.tanasescu@nxp.com>
Date:   Tue May 30 17:21:59 2023 +0300

    Bluetooth: ISO: Add support for connecting multiple BISes

    [ Upstream commit a0bfde167b506423111ddb8cd71930497a40fc54 ]

    It is required for some configurations to have multiple BISes as part
    of the same BIG.

    Similar to the flow implemented for unicast, DEFER_SETUP will also be
    used to bind multiple BISes for the same BIG, before starting Periodic
    Advertising and creating the BIG.

    The user will have to open a new socket for each BIS. By setting the
    BT_DEFER_SETUP socket option and calling connect, a new connection
    will be added for the BIG and advertising handle set by the socket
    QoS parameters. Since all BISes will be bound for the same BIG and
    advertising handle, the socket QoS options and base parameters should
    match for all connections.
```

```
    By calling connect on a socket that does not have the BT_DEFER_SETUP
    option set, periodic advertising will be started and the BIG will
    be created, with a BIS for each previously bound connection. Since
    a BIG cannot be reconfigured with additional BISes after creation,
    no more connections can be bound for the BIG after the start periodic
    advertising and create BIG commands have been queued.

    The bis_cleanup function has also been updated, so that the advertising
    set and the BIG will not be terminated unless there are no more
    bound or connected BISes.

    The HCI_CONN_BIG_CREATED connection flag has been added to indicate
    that the BIG has been successfully created. This flag is checked at
    bis_cleanup, so that the BIG is only terminated if the
    HCI_LE_Create_BIG_Complete has been received.

    This implementation has been tested on hardware, using the "isotest"
    tool with an additional command line option, to specify the number of
    BISes to create as part of the desired BIG:

        tools/isotest -i hci0 -s 00:00:00:00:00:00 -N 2 -G 1 -T 1

    The btmon log shows that a BIG containing 2 BISes has been created:

    < HCI Command: LE Create Broadcast Isochronous Group (0x08|0x0068) plen 31
            Handle: 0x01
            Advertising Handle: 0x01
            Number of BIS: 2
            SDU Interval: 10000 us (0x002710)
            Maximum SDU size: 40
            Maximum Latency: 10 ms (0x000a)
            RTN: 0x02
            PHY: LE 2M (0x02)
            Packing: Sequential (0x00)
            Framing: Unframed (0x00)
            Encryption: 0x00
            Broadcast Code: 00000000000000000000000000000000

    > HCI Event: Command Status (0x0f) plen 4
          LE Create Broadcast Isochronous Group (0x08|0x0068) ncmd 1
            Status: Success (0x00)

    > HCI Event: LE Meta Event (0x3e) plen 23
          LE Broadcast Isochronous Group Complete (0x1b)
            Status: Success (0x00)
            Handle: 0x01
            BIG Synchronization Delay: 1974 us (0x0007b6)
            Transport Latency: 1974 us (0x0007b6)
            PHY: LE 2M (0x02)
            NSE: 3
            BN: 1
            PTO: 1
            IRC: 3
            Maximum PDU: 40
            ISO Interval: 10.00 msec (0x0008)
            Connection Handle #0: 10
            Connection Handle #1: 11

    < HCI Command: LE Setup Isochronous Data Path (0x08|0x006e) plen 13
            Handle: 10
            Data Path Direction: Input (Host to Controller) (0x00)
            Data Path: HCI (0x00)
            Coding Format: Transparent (0x03)
            Company Codec ID: Ericsson Technology Licensing (0)
```

```
              Vendor Codec ID: 0
              Controller Delay: 0 us (0x000000)
              Codec Configuration Length: 0
              Codec Configuration:

      > HCI Event: Command Complete (0x0e) plen 6
            LE Setup Isochronous Data Path (0x08|0x006e) ncmd 1
              Status: Success (0x00)
              Handle: 10

      < HCI Command: LE Setup Isochronous Data Path (0x08|0x006e) plen 13
              Handle: 11
              Data Path Direction: Input (Host to Controller) (0x00)
              Data Path: HCI (0x00)
              Coding Format: Transparent (0x03)
              Company Codec ID: Ericsson Technology Licensing (0)
              Vendor Codec ID: 0
              Controller Delay: 0 us (0x000000)
              Codec Configuration Length: 0
              Codec Configuration:

      > HCI Event: Command Complete (0x0e) plen 6
            LE Setup Isochronous Data Path (0x08|0x006e) ncmd 1
              Status: Success (0x00)
              Handle: 11

      < ISO Data TX: Handle 10 flags 0x02 dlen 44

      < ISO Data TX: Handle 11 flags 0x02 dlen 44

      > HCI Event: Number of Completed Packets (0x13) plen 5
              Num handles: 1
              Handle: 10
              Count: 1

      > HCI Event: Number of Completed Packets (0x13) plen 5
              Num handles: 1
              Handle: 11
              Count: 1

      Signed-off-by: Iulia Tanasescu <iulia.tanasescu@nxp.com>
      Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
      Stable-dep-of: 7f74563e6140 ("Bluetooth: ISO: do not emit new LE Create CIS if previous
  is pending")
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit c0dbcebc7f390ec7dbe010dcc22c60f0c6bfc26d
  Author: Herbert Xu <herbert@gondor.apana.org.au>
  Date:   Thu Aug 3 17:59:28 2023 +0800

      crypto: api - Use work queue in crypto_destroy_instance

      [ Upstream commit 9ae4577bc077a7e32c3c7d442c95bc76865c0f17 ]

      The function crypto_drop_spawn expects to be called in process
      context.  However, when an instance is unregistered while it still
      has active users, the last user may cause the instance to be freed
      in atomic context.

      Fix this by delaying the freeing to a work queue.

      Fixes: 6bfd48096ff8 ("[CRYPTO] api: Added spawns")
      Reported-by: Florent Revest <revest@chromium.org>
      Reported-by: syzbot+d769eed29cc42d75e2a3@syzkaller.appspotmail.com
      Reported-by: syzbot+610ec0671f51e838436e@syzkaller.appspotmail.com
```

```
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Tested-by: Florent Revest <revest@chromium.org>
    Acked-by: Florent Revest <revest@chromium.org>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2de34b25a0ab549d7f0c6218ac90109e0a5e872a
Author: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
Date:   Mon Jul 31 18:54:54 2023 +0200

    crypto: stm32 - Properly handle pm_runtime_get failing

    [ Upstream commit aec48805163338f8413118796c1dd035661b9140 ]

    If pm_runtime_get() (disguised as pm_runtime_resume_and_get()) fails, this
    means the clk wasn't prepared and enabled. Returning early in this case
    however is wrong as then the following resource frees are skipped and this
    is never catched up. So do all the cleanups but clk_disable_unprepare().

    Also don't emit a warning, as stm32_hash_runtime_resume() already emitted
    one.

    Note that the return value of stm32_hash_remove() is mostly ignored by
    the device core. The only effect of returning zero instead of an error
    value is to suppress another warning in platform_remove(). So return 0
    even if pm_runtime_resume_and_get() failed.

    Fixes: 8b4d566de6a5 ("crypto: stm32/hash - Add power management support")
    Signed-off-by: Uwe Kleine-König <u.kleine-koenig@pengutronix.de>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 93482c1a6ca1b0664c09566df1fdd6fa64383b5a
Author: Miguel Ojeda <ojeda@kernel.org>
Date:   Fri Jun 16 02:16:27 2023 +0200

    kbuild: rust_is_available: fix confusion when a version appears in the path

    [ Upstream commit 9eb7e20e0c5cd069457845f965b3e8a7d736ecb7 ]

    `bindgen`'s output for `libclang`'s version check contains paths, which
    in turn may contain strings that look like version numbers [1][2]:

        .../6.1.0-dev/.../rust_is_available_bindgen_libclang.h:2:9: warning: clang version
11.1.0  [-W#pragma-messages], err: false

    which the script will pick up as the version instead of the latter.

    It is also the case that versions may appear after the actual version
    (e.g. distribution's version text), which was the reason behind `head` [3]:

        .../rust-is-available-bindgen-libclang.h:2:9: warning: clang version 13.0.0 (Fedora
13.0.0-3.fc35) [-W#pragma-messages], err: false

    Thus instead ask for a match after the `clang version` string.

    Reported-by: Jordan Isaacs <mail@jdisaacs.com>
    Closes: https://github.com/Rust-for-Linux/linux/issues/942 [1]
    Reported-by: "Ethan D. Twardy" <ethan.twardy@gmail.com>
    Closes: https://lore.kernel.org/rust-for-linux/20230528131802.6390-2-
ethan.twardy@gmail.com/ [2]
    Reported-by: Tiago Lam <tiagolam@gmail.com>
    Closes: https://github.com/Rust-for-Linux/linux/pull/789 [3]
    Fixes: 78521f3399ab ("scripts: add `rust_is_available.sh`")
    Reviewed-by: Martin Rodriguez Reboredo <yakoyoku@gmail.com>
```

```
    Reviewed-by: Ethan Twardy <ethan.twardy@gmail.com>
    Tested-by: Ethan Twardy <ethan.twardy@gmail.com>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20230616001631.463536-8-ojeda@kernel.org
    Signed-off-by: Miguel Ojeda <ojeda@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 337ff828ba1e33dd1ce75d61c60465ea2182f135
Author: Miguel Ojeda <ojeda@kernel.org>
Date:   Fri Jun 16 02:16:25 2023 +0200

    kbuild: rust_is_available: add check for `bindgen` invocation

    [ Upstream commit 52cae7f28ed6c3992489f16bb355f5b623f0912e ]

    `scripts/rust_is_available.sh` calls `bindgen` with a special
    header in order to check whether the `libclang` version in use
    is suitable.

    However, the invocation itself may fail if, for instance, `bindgen`
    cannot locate `libclang`. This is fine for Kconfig (since the
    script will still fail and therefore disable Rust as it should),
    but it is pretty confusing for users of the `rustavailable` target
    given the error will be unrelated:

        ./scripts/rust_is_available.sh: 21: arithmetic expression: expecting primary:
"100000 *  + 100 *  + "
        make: *** [Makefile:1816: rustavailable] Error 2

    Instead, run the `bindgen` invocation independently in a previous
    step, saving its output and return code. If it fails, then show
    the user a proper error message. Otherwise, continue as usual
    with the saved output.

    Since the previous patch we show a reference to the docs, and
    the docs now explain how `bindgen` looks for `libclang`,
    thus the error message can leverage the documentation, avoiding
    duplication here (and making users aware of the setup guide in
    the documentation).

    Reported-by: Nick Desaulniers <ndesaulniers@google.com>
    Link: https://lore.kernel.org/rust-for-
linux/CAKwvOdm5JT4wbdQQYuW+RT07rCi6whGBM2iUAyg8A1CmLXG6Nw@mail.gmail.com/
    Reported-by: François Valenduc <francoisvalenduc@gmail.com>
    Closes: https://github.com/Rust-for-Linux/linux/issues/934
    Reported-by: Alexandru Radovici <msg4alex@gmail.com>
    Closes: https://github.com/Rust-for-Linux/linux/pull/921
    Reported-by: Matthew Leach <dev@mattleach.net>
    Closes: https://lore.kernel.org/rust-for-linux/20230507084116.1099067-1-
dev@mattleach.net/
    Fixes: 78521f3399ab ("scripts: add `rust_is_available.sh`")
    Reviewed-by: Martin Rodriguez Reboredo <yakoyoku@gmail.com>
    Reviewed-by: Masahiro Yamada <masahiroy@kernel.org>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20230616001631.463536-6-ojeda@kernel.org
    Signed-off-by: Miguel Ojeda <ojeda@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 36cd07efa2b56b21d52d54ad08a1af2cecbf5481
Author: Russell Currey <ruscur@russell.cc>
Date:   Fri Jun 16 02:16:22 2023 +0200

    kbuild: rust_is_available: fix version check when CC has multiple arguments

    [ Upstream commit dee3a6b819c96fc8b1907577f585fd66f5c0fefe ]
```

```
    rust_is_available.sh uses cc-version.sh to identify which C compiler is
    in use, as scripts/Kconfig.include does.  cc-version.sh isn't designed to
    be able to handle multiple arguments in one variable, i.e. "ccache clang".
    Its invocation in rust_is_available.sh quotes "$CC", which makes
    $1 == "ccache clang" instead of the intended $1 == ccache & $2 == clang.

    cc-version.sh could also be changed to handle having "ccache clang" as one
    argument, but it only has the one consumer upstream, making it simpler to
    fix the caller here.

    Signed-off-by: Russell Currey <ruscur@russell.cc>
    Fixes: 78521f3399ab ("scripts: add `rust_is_available.sh`")
    Link: https://github.com/Rust-for-Linux/linux/pull/873
    [ Reworded title prefix and reflow line to 75 columns. ]
    Reviewed-by: Martin Rodriguez Reboredo <yakoyoku@gmail.com>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20230616001631.463536-3-ojeda@kernel.org
    Signed-off-by: Miguel Ojeda <ojeda@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 237a02a2cda1f25211c302af24827fc97ba8c273
Author: Masahiro Yamada <masahiroy@kernel.org>
Date:   Fri Jun 16 02:16:21 2023 +0200

    kbuild: rust_is_available: remove -v option

    [ Upstream commit d824d2f98565e7c4cb1b862c230198fbe1a968be ]

    The -v option is passed when this script is invoked from Makefile,
    but not when invoked from Kconfig.

    As you can see in scripts/Kconfig.include, the 'success' macro suppresses
    stdout and stderr anyway, so this script does not need to be quiet.

    Signed-off-by: Masahiro Yamada <masahiroy@kernel.org>
    Reviewed-by: Miguel Ojeda <ojeda@kernel.org>
    Tested-by: Miguel Ojeda <ojeda@kernel.org>
    Reviewed-by: Nathan Chancellor <nathan@kernel.org>
    Link: https://lore.kernel.org/r/20230109061436.3146442-1-masahiroy@kernel.org
    [ Reworded prefix to match the others in the patch series. ]
    Reviewed-by: Martin Rodriguez Reboredo <yakoyoku@gmail.com>
    Link: https://lore.kernel.org/r/20230616001631.463536-2-ojeda@kernel.org
    Signed-off-by: Miguel Ojeda <ojeda@kernel.org>
    Stable-dep-of: dee3a6b819c9 ("kbuild: rust_is_available: fix version check when CC has
multiple arguments")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b13e18d53012452645d0a0f853277ce91023bb45
Author: Kui-Feng Lee <thinker.li@gmail.com>
Date:   Thu Aug 3 16:12:06 2023 -0700

    bpf: fix bpf_dynptr_slice() to stop return an ERR_PTR.

    [ Upstream commit 5426700e6841bf72e652e34b5cec68eadf442435 ]

    Verify if the pointer obtained from bpf_xdp_pointer() is either an error or
    NULL before returning it.

    The function bpf_dynptr_slice() mistakenly returned an ERR_PTR. Instead of
    solely checking for NULL, it should also verify if the pointer returned by
    bpf_xdp_pointer() is an error or NULL.

    Reported-by: Dan Carpenter <dan.carpenter@linaro.org>
    Closes: https://lore.kernel.org/bpf/d1360219-85c3-4a03-9449-
```

```
 253ea905f9d1@moroto.mountain/
     Fixes: 66e3a13e7c2c ("bpf: Add bpf_dynptr_slice and bpf_dynptr_slice_rdwr")
     Suggested-by: Alexei Starovoitov <alexei.starovoitov@gmail.com>
     Signed-off-by: Kui-Feng Lee <thinker.li@gmail.com>
     Acked-by: Yonghong Song <yonghong.song@linux.dev>
     Link: https://lore.kernel.org/r/20230803231206.1060485-1-thinker.li@gmail.com
     Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 2c9d205040d7c0eaccc473917f9b0bb0a923e440
 Author: David Howells <dhowells@redhat.com>
 Date:   Wed Jul 26 22:53:19 2023 +0100

     crypto: af_alg - Fix missing initialisation affecting gcm-aes-s390

     [ Upstream commit 6a4b8aa0a916b39a39175584c07222434fa6c6ef ]

     Fix af_alg_alloc_areq() to initialise areq->first_rsgl.sgl.sgt.sgl to point
     to the scatterlist array in areq->first_rsgl.sgl.sgl.

     Without this, the gcm-aes-s390 driver will oops when it tries to do
     gcm_walk_start() on req->dst because req->dst is set to the value of
     areq->first_rsgl.sgl.sgl by _aead_recvmsg() calling
     aead_request_set_crypt().

     The problem comes if an empty ciphertext is passed: the loop in
     af_alg_get_rsgl() just passes straight out and doesn't set areq->first_rsgl
     up.

     This isn't a problem on x86_64 using gcmaes_crypt_by_sg() because, as far
     as I can tell, that ignores req->dst and only uses req->src[*].

     [*] Is this a bug in aesni-intel_glue.c?

     The s390x oops looks something like:

      Unable to handle kernel pointer dereference in virtual kernel address space
      Failing address: 0000000a00000000 TEID: 0000000a00000803
      Fault in home space mode while using kernel ASCE.
      AS:00000000a43a0007 R3:0000000000000024
      Oops: 003b ilc:2 [#1] SMP
      ...
      Call Trace:
       [<000003ff7fc3d47e>] gcm_walk_start+0x16/0x28 [aes_s390]
       [<00000000a2a342f2>] crypto_aead_decrypt+0x9a/0xb8
       [<00000000a2a60888>] aead_recvmsg+0x478/0x698
       [<00000000a2e519a0>] sock_recvmsg+0x70/0xb0
       [<00000000a2e51a56>] sock_read_iter+0x76/0xa0
       [<00000000a273e066>] vfs_read+0x26e/0x2a8
       [<00000000a273e8c4>] ksys_read+0xbc/0x100
       [<00000000a311d808>] __do_syscall+0x1d0/0x1f8
       [<00000000a312ff30>] system_call+0x70/0x98
      Last Breaking-Event-Address:
       [<000003ff7fc3e6b4>] gcm_aes_crypt+0x104/0xa68 [aes_s390]

     Fixes: c1abe6f570af ("crypto: af_alg: Use extract_iter_to_sg() to create scatterlists")
     Reported-by: Ondrej Mosnáček <omosnacek@gmail.com>
     Link:
 https://lore.kernel.org/r/CAAUqJDuRkHE8fPgZJGaKjUjd3QfGwzfumuJBmStPqBhubxyk_A@mail.gmail.com
 /
     Signed-off-by: David Howells <dhowells@redhat.com>
     cc: Herbert Xu <herbert@gondor.apana.org.au>
     cc: Sven Schnelle <svens@linux.ibm.com>
     cc: Harald Freudenberger <freude@linux.vnet.ibm.com>
     cc: "David S. Miller" <davem@davemloft.net>
```

```
    cc: Paolo Abeni <pabeni@redhat.com>
    cc: linux-crypto@vger.kernel.org
    cc: linux-s390@vger.kernel.org
    cc: regressions@lists.linux.dev
    Tested-by: Sven Schnelle <svens@linux.ibm.com>
    Tested-by: Ondrej Mosnáček <omosnacek@gmail.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f24d6834a3685bb5214fa0dd34f4d6eae61c6fd8
Author: Li Zetao <lizetao1@huawei.com>
Date:   Thu Aug 3 21:48:05 2023 +0800

    spi: mpc5xxx-psc: Fix unsigned expression compared with zero

    [ Upstream commit de5e92cb5cefd2968b96075995a36e28298edf71 ]

    There is two warnings reported by coccinelle:

    ./drivers/spi/spi-mpc512x-psc.c:493:5-13: WARNING:
            Unsigned expression compared with zero: mps -> irq     <     0
    ./drivers/spi/spi-mpc52xx-psc.c:332:5-13: WARNING:
            Unsigned expression compared with zero: mps -> irq     <     0

    The commit "208ee586f862"
    ("spi: mpc5xxx-psc: Return immediately if IRQ resource is unavailable")
    was to check whether the IRQ resource is unavailable. When the IRQ
    resource is unavailable, an error code is returned, however, the type
    of "mps->irq" is "unsigned int", causing the error code to flip. Modify
    the type of "mps->irq" to solve this problem.

    Fixes: 208ee586f862 ("spi: mpc5xxx-psc: Return immediately if IRQ resource is
unavailable")
    Signed-off-by: Li Zetao <lizetao1@huawei.com>
    Link: https://lore.kernel.org/r/20230803134805.1037251-1-lizetao1@huawei.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 44cc7f44dd178b3edfec2d2ecd2d88f014a0cc30
Author: Yu Liao <liaoyu15@huawei.com>
Date:   Wed Aug 2 10:31:30 2023 +0800

    power: supply: qcom_pmi8998_charger: fix uninitialized variable

    [ Upstream commit 13a0d1088c8fea1565e30a169188b59bdd77759e ]

    smatch warnings:
      drivers/power/supply/qcom_pmi8998_charger.c:565 smb2_status_change_work() error:
uninitialized symbol 'usb_online'.

    usb_online is used uninitialized whenever smb2_get_prop_usb_online()
    returns a negative value.

    Thus, fix the issue by initializing usb_online to 0.

    Fixes: 8648aeb5d7b7 ("power: supply: add Qualcomm PMI8998 SMB2 Charger driver")
    Reported-by: kernel test robot <lkp@intel.com>
    Closes: https://lore.kernel.org/r/202307280638.556PrzIS-lkp@intel.com/
    Signed-off-by: Yu Liao <liaoyu15@huawei.com>
    Reviewed-by: Caleb Connolly <caleb.connolly@linaro.org
    Link: https://lore.kernel.org/r/20230802023130.2516232-1-liaoyu15@huawei.com
    Signed-off-by: Sebastian Reichel <sebastian.reichel@collabora.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c0e5be559a0637a3e01a653d5d6794e1dbf8df5c
```

```
Author: Alan Maguire <alan.maguire@oracle.com>
Date:   Wed Aug 2 08:39:06 2023 +0100

    selftests/bpf: fix static assert compilation issue for test_cls_*.c

    [ Upstream commit 416c6d01244ecbf0abfdb898fd091b50ef951b48 ]

    commit bdeeed3498c7 ("libbpf: fix offsetof() and container_of() to work with CO-RE")

    ...was backported to stable trees such as 5.15. The problem is that with older
    LLVM/clang (14/15) - which is often used for older kernels - we see compilation
    failures in BPF selftests now:

    In file included from progs/test_cls_redirect_subprogs.c:2:
    progs/test_cls_redirect.c:90:2: error: static assertion expression is not an integral
constant expression
            sizeof(flow_ports_t) !=
            ^~~~~~~~~~~~~~~~~~~~~
    progs/test_cls_redirect.c:91:3: note: cast that performs the conversions of a
reinterpret_cast is not allowed in a constant expression
                    offsetofend(struct bpf_sock_tuple, ipv4.dport) -
                    ^
    progs/test_cls_redirect.c:32:3: note: expanded from macro 'offsetofend'
            (offsetof(TYPE, MEMBER) + sizeof((((TYPE *)0)->MEMBER)))
             ^
    tools/testing/selftests/bpf/tools/include/bpf/bpf_helpers.h:86:33: note: expanded from
macro 'offsetof'
                                    ^
    In file included from progs/test_cls_redirect_subprogs.c:2:
    progs/test_cls_redirect.c:95:2: error: static assertion expression is not an integral
constant expression
            sizeof(flow_ports_t) !=
            ^~~~~~~~~~~~~~~~~~~~~
    progs/test_cls_redirect.c:96:3: note: cast that performs the conversions of a
reinterpret_cast is not allowed in a constant expression
                    offsetofend(struct bpf_sock_tuple, ipv6.dport) -
                    ^
    progs/test_cls_redirect.c:32:3: note: expanded from macro 'offsetofend'
            (offsetof(TYPE, MEMBER) + sizeof((((TYPE *)0)->MEMBER)))
             ^
    tools/testing/selftests/bpf/tools/include/bpf/bpf_helpers.h:86:33: note: expanded from
macro 'offsetof'
                                    ^
    2 errors generated.
    make: *** [Makefile:594: tools/testing/selftests/bpf/test_cls_redirect_subprogs.bpf.o]
Error 1

    The problem is the new offsetof() does not play nice with static asserts.
    Given that the context is a static assert (and CO-RE relocation is not
    needed at compile time), offsetof() usage can be replaced by restoring
    the original offsetof() definition as __builtin_offsetof().

    Fixes: bdeeed3498c7 ("libbpf: fix offsetof() and container_of() to work with CO-RE")
    Reported-by: Colm Harrington <colm.harrington@oracle.com>
    Signed-off-by: Alan Maguire <alan.maguire@oracle.com>
    Tested-by: Yipeng Zou <zouyipeng@huawei.com>
    Acked-by: Yonghong Song <yonghong.song@linux.dev>
    Link: https://lore.kernel.org/r/20230802073906.3197480-1-alan.maguire@oracle.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0d622342c26c20adb5f0315333e99d67f6cf494d
Author: Arnd Bergmann <arnd@arndb.de>
Date:   Mon Jul 3 14:37:29 2023 +0200
```

```
    wifi: ath12k: fix memcpy array overflow in ath12k_peer_assoc_h_he()

    [ Upstream commit 603cf6c2fcdcbc38f1daa316794e7268852677a7 ]

    Two memory copies in this function copy from a short array into a longer one,
    using the wrong size, which leads to an out-of-bounds access:

    include/linux/fortify-string.h:592:4: error: call to '__read_overflow2_field' declared
with 'warning' attribute: detected read beyond size of field (2nd parameter); maybe use
struct_group()? [-Werror,-Wattribute-warning]
                            __read_overflow2_field(q_size_field, size);
                            ^
    include/linux/fortify-string.h:592:4: error: call to '__read_overflow2_field' declared
with 'warning' attribute: detected read beyond size of field (2nd parameter); maybe use
struct_group()? [-Werror,-Wattribute-warning]
    2 errors generated.

    Fixes: d889913205cf7 ("wifi: ath12k: driver for Qualcomm Wi-Fi 7 devices")
    Signed-off-by: Arnd Bergmann <arnd@arndb.de>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230703123737.3420464-1-arnd@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7fe3cde7223f9449010fb1163e6a26ecdfcca9b1
Author: Aditya Kumar Singh <quic_adisi@quicinc.com>
Date:   Wed Jul 26 10:16:24 2023 +0530

    wifi: ath11k: fix band selection for ppdu received in channel 177 of 5 GHz

    [ Upstream commit 72c8caf904aed2caed5d6e75233294b6159ddb5d ]

    5 GHz band channel 177 support was added with the commit e5e94d10c856 ("wifi:
    ath11k: add channel 177 into 5 GHz channel list"). However, during processing
    for the received ppdu in ath11k_dp_rx_h_ppdu(), channel number is checked only
    till 173. This leads to driver code checking for channel and then fetching the
    band from it which is extra effort since firmware has already given the channel
    number in the metadata.

    Fix this issue by checking the channel number till 177 since we support
    it now.

    Found via code review. Compile tested only.

    Fixes: e5e94d10c856 ("wifi: ath11k: add channel 177 into 5 GHz channel list")
    Signed-off-by: Aditya Kumar Singh <quic_adisi@quicinc.com>
    Signed-off-by: Kalle Valo <quic_kvalo@quicinc.com>
    Link: https://lore.kernel.org/r/20230726044624.20507-1-quic_adisi@quicinc.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7d13b9048d20a799179bd16ee9f38ed15ea88289
Author: Dmitry Antipov <dmantipov@yandex.ru>
Date:   Mon Jul 31 10:43:07 2023 +0300

    wifi: mwifiex: fix error recovery in PCIE buffer descriptor management

    [ Upstream commit 288c63d5cb4667a51a04668b3e2bb0ea499bc5f4 ]

    Add missing 'kfree_skb()' in 'mwifiex_init_rxq_ring()' and never do
    'kfree(card->rxbd_ring_vbase)' because this area is DMAed and should
    be released with 'dma_free_coherent()'. The latter is performed in
    'mwifiex_pcie_delete_rxbd_ring()', which is now called to recover
    from possible errors in 'mwifiex_pcie_create_rxbd_ring()'. Likewise
    for 'mwifiex_pcie_init_evt_ring()', 'kfree(card->evtbd_ring_vbase)'
    'mwifiex_pcie_delete_evtbd_ring()' and 'mwifiex_pcie_create_rxbd_ring()'.
```

```
    Fixes: d930faee141b ("mwifiex: add support for Marvell pcie8766 chipset")
    Signed-off-by: Dmitry Antipov <dmantipov@yandex.ru>
    Acked-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230731074334.56463-1-dmantipov@yandex.ru
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 650d1bc02fba7b42f476d8b6643324abac5921ed
Author: Polaris Pi <pinkperfect2021@gmail.com>
Date:   Sun Jul 23 07:07:41 2023 +0000

    wifi: mwifiex: Fix OOB and integer underflow when rx packets

    [ Upstream commit 11958528161731c58e105b501ed60b83a91ea941 ]

    Make sure mwifiex_process_mgmt_packet,
    mwifiex_process_sta_rx_packet and mwifiex_process_uap_rx_packet,
    mwifiex_uap_queue_bridged_pkt and mwifiex_process_rx_packet
    not out-of-bounds access the skb->data buffer.

    Fixes: 2dbaf751b1de ("mwifiex: report received management frames to cfg80211")
    Signed-off-by: Polaris Pi <pinkperfect2021@gmail.com>
    Reviewed-by: Matthew Wang <matthewmwang@chromium.org>
    Reviewed-by: Brian Norris <briannorris@chromium.org>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230723070741.1544662-1-pinkperfect2021@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 37d9b131129302206798b1fd44d85d66921c29dd
Author: Larry Finger <Larry.Finger@lwfinger.net>
Date:   Mon Jul 24 13:39:27 2023 -0500

    wifi: rtw89: Fix loading of compressed firmware

    [ Upstream commit 942999c48cb382feb53c6da7679a994c97963836 ]

    When using compressed firmware, the early firmware load feature will fail.
    In most cases, the only downside is that if a device has more than one
    firmware version available, only the last one listed will be loaded.
    In at least two cases, there is no firmware loaded, and the device fails
    initialization. See https://github.com/lwfinger/rtw89/issues/259 and
    https://bugzilla.opensuse.org/show_bug.cgi?id=1212808 for examples of
    the failure.

    When firmware_class.dyndbg=+p" added to the kernel boot parameters, the
    following is found:

    finger@localhost:~/rtw89>sudo dmesg -t | grep rtw89
    firmware_class: __allocate_fw_priv: fw-rtw89/rtw8852b_fw-1.bin fw_priv=00000000638862fb
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/5.14.21-150500.53-
default/rtw89/rtw8852b_fw-1.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw-1.bin failed
for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/5.14.21-150500.53-
default/rtw89/rtw8852b_fw-1.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/rtw89/rtw8852b_fw-1.bin failed for no
such file or directory.
    rtw89_8852be 0000:02:00.0: Direct firmware load for rtw89/rtw8852b_fw-1.bin failed with
error -2
    firmware_class: __free_fw_priv: fw-rtw89/rtw8852b_fw-1.bin fw_priv=00000000638862fb
data=00000000307c30c7 size=0
    firmware_class: __allocate_fw_priv: fw-rtw89/rtw8852b_fw.bin fw_priv=00000000638862fb
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw.bin failed
```

```
for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/rtw89/rtw8852b_fw.bin failed for no
such file or directory.
    rtw89_8852be 0000:02:00.0: Direct firmware load for rtw89/rtw8852b_fw.bin failed with
error -2
    firmware_class: __free_fw_priv: fw-rtw89/rtw8852b_fw.bin fw_priv=00000000638862fb
data=00000000307c30c7 size=0
    rtw89_8852be 0000:02:00.0: failed to early request firmware: -2
    firmware_class: __allocate_fw_priv: fw-rtw89/rtw8852b_fw.bin fw_priv=00000000638862fb
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw.bin failed
for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/rtw89/rtw8852b_fw.bin failed for no
such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin.xz failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw.bin.xz failed
for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/5.14.21-150500.53-
default/rtw89/rtw8852b_fw.bin.xz failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: Loading firmware from /lib/firmware/rtw89/rtw8852b_fw.bin.xz
    rtw89_8852be 0000:02:00.0: f/w decompressing rtw89/rtw8852b_fw.bin
    firmware_class: fw_set_page_data: fw-rtw89/rtw8852b_fw.bin fw_priv=00000000638862fb
data=000000004ed6c2f7 size=1035232
    rtw89_8852be 0000:02:00.0: Firmware version 0.27.32.1, cmd version 0, type 1
    rtw89_8852be 0000:02:00.0: Firmware version 0.27.32.1, cmd version 0, type 3

    The key is that firmware version 0.27.32.1 is loaded.

    With this patch, the following is obtained:

    firmware_class: __free_fw_priv: fw-rtw89/rtw8852b_fw.bin fw_priv=000000000849addc
data=00000000fd3cabe2 size=1035232
    firmware_class: fw_name_devm_release: fw_name-rtw89/rtw8852b_fw.bin devm-
000000002d8c3343 released
    firmware_class: __allocate_fw_priv: fw-rtw89/rtw8852b_fw-1.bin fw_priv=000000009e1a6364
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/6.4.3-1-
default/rtw89/rtw8852b_fw-1.bin failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw-1.bin failed
for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/6.4.3-1-default/rtw89/rtw8852b_fw-1.bin
failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/rtw89/rtw8852b_fw-1.bin failed for no
such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/6.4.3-1-
default/rtw89/rtw8852b_fw-1.bin.zst failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw-1.bin.zst
failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/6.4.3-1-default/rtw89/rtw8852b_fw-
1.bin.zst failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/rtw89/rtw8852b_fw-1.bin.zst failed for
no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/6.4.3-1-
default/rtw89/rtw8852b_fw-1.bin.xz failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/updates/rtw89/rtw8852b_fw-1.bin.xz
failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: loading /lib/firmware/6.4.3-1-default/rtw89/rtw8852b_fw-
1.bin.xz failed for no such file or directory.
    rtw89_8852be 0000:02:00.0: Loading firmware from /lib/firmware/rtw89/rtw8852b_fw-
1.bin.xz
```

```
    rtw89_8852be 0000:02:00.0: f/w decompressing rtw89/rtw8852b_fw-1.bin
    firmware_class: fw_set_page_data: fw-rtw89/rtw8852b_fw-1.bin fw_priv=000000009e1a6364
 data=00000000fd3cabe2 size=1184992
    rtw89_8852be 0000:02:00.0: Loaded FW: rtw89/rtw8852b_fw-1.bin, sha256:
 8539efc75f513f4585cf0cd6e79e6507da47fce87225f2d0de391a03aefe9ac8
    rtw89_8852be 0000:02:00.0: loaded firmware rtw89/rtw8852b_fw-1.bin
    rtw89_8852be 0000:02:00.0: Firmware version 0.29.29.1, cmd version 0, type 5
    rtw89_8852be 0000:02:00.0: Firmware version 0.29.29.1, cmd version 0, type 3

    Now, version 0.29.29.1 is loaded.

    Fixes: ffde7f3476a6 ("wifi: rtw89: add firmware format version to backward compatible
 with older drivers")
    Cc: Ping-Ke Shih <pkshih@realtek.com>
    Cc: Takashi Iwai <tiwai@suse.de>
    Signed-off-by: Larry Finger <Larry.Finger@lwfinger.net>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/20230724183927.28553-1-Larry.Finger@lwfinger.net
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c43017fbebcc365881509f767a550bf9f54df6d3
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Thu Jul 27 02:35:06 2023 +0800

    wifi: mt76: mt7915: fix power-limits while chan_switch

    [ Upstream commit 6c0570bc21ec2073890aa252c8420ca7bec402e4 ]

    If user changes the channel without completely disabling the interface the
    txpower_sku values reported track the old channel the device was operating on.
    If user bounces the interface the correct power tables are applied.

    mt7915_sku_group_len array gets updated before the channel switch happens so it
    uses data from the old channel.

    Fixes: ecb187a74e18 ("mt76: mt7915: rework the flow of txpower setting")
    Fixes: f1d962369d56 ("mt76: mt7915: implement HE per-rate tx power support")
    Reported-By: Chad Monroe <chad.monroe@smartrg.com>
    Tested-by: Chad Monroe <chad.monroe@smartrg.com>
    Signed-off-by: Allen Ye <allen.ye@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit edb1afe042c740d97cd7931f3734e82d5397ad4f
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Thu Jul 27 02:35:05 2023 +0800

    wifi: mt76: mt7915: fix tlv length of mt7915_mcu_get_chan_mib_info

    [ Upstream commit 4f1875c288dfc1ccea81fc17fef1d30c9d8498b2 ]

    Correct per-device TLV lengths to avoid invalid operation in firmware.
    (  64.040375:28:STATS-E)statsGetSingleHWCounter: MIB counter index = 65472 not
 supported.
    This happens on mt7916/mt7986.

    Fixes: b0bfa00595be ("wifi: mt76: mt7915: improve accuracy of time_busy calculation")
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 04fc9843d181814f3d6336723f519114167f1d01
Author: Markus Schneider-Pargmann <msp@baylibre.com>
Date:   Fri Jul 28 16:19:19 2023 +0200
```

```
    can: tcan4x5x: Remove reserved register 0x814 from writable table

    [ Upstream commit fbe534f7bf213d485b0ed5362b24a41bf3e18803 ]

    The mentioned register is not writable. It is reserved and should not be
    written.

    Fixes: 39dbb21b6a29 ("can: tcan4x5x: Specify separate read/write ranges")
    Signed-off-by: Markus Schneider-Pargmann <msp@baylibre.com>
    Reviewed-by: Michal Kubiak <michal.kubiak@intel.com>
    Link: https://lore.kernel.org/all/20230728141923.162477-3-msp@baylibre.com
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03cd9a222e1cca710e12ee179591a66ca59ee4c7
Author: Marc Kleine-Budde <mkl@pengutronix.de>
Date:   Tue Jul 4 11:23:37 2023 +0200

    can: gs_usb: gs_usb_receive_bulk_callback(): count RX overflow errors also in case of
OOM

    [ Upstream commit 6c8bc15f02b85bc8f47074110d8fd8caf7a1e42d ]

    In case of an RX overflow error from the CAN controller and an OOM
    where no skb can be allocated, the error counters are not incremented.

    Fix this by first incrementing the error counters and then allocate
    the skb.

    Fixes: d08e973a77d1 ("can: gs_usb: Added support for the GS_USB CAN devices")
    Link: https://lore.kernel.org/all/20230718-gs_usb-cleanups-v1-7-
c3b9154ec605@pengutronix.de
    Signed-off-by: Marc Kleine-Budde <mkl@pengutronix.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 41e2d4e0210bb531ac60f7fb1e9bc1211fdd69b9
Author: Zhang Shurong <zhang_shurong@foxmail.com>
Date:   Sat Jul 22 23:49:09 2023 +0800

    spi: tegra20-sflash: fix to check return value of platform_get_irq() in
tegra_sflash_probe()

    [ Upstream commit 29a449e765ff70a5bd533be94babb6d36985d096 ]

    The platform_get_irq might be failed and return a negative result. So
    there should have an error handling code.

    Fixed this by adding an error handling code.

    Fixes: 8528547bcc33 ("spi: tegra: add spi driver for sflash controller")
    Signed-off-by: Zhang Shurong <zhang_shurong@foxmail.com>
    Link: https://lore.kernel.org/r/tencent_71FC162D589E4788C2152AAC84CD8D5C6D06@qq.com
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3e4f7dedaec9ef3c9567c6e0a65412e59efad314
Author: Lin Ma <linma@zju.edu.cn>
Date:   Sun Jul 23 16:03:50 2023 +0800

    wifi: mt76: testmode: add nla_policy for MT76_TM_ATTR_TX_LENGTH

    [ Upstream commit 74f12d511625e603fac8c0c2b6872e687e56dd61 ]

    It seems that the nla_policy in mt76_tm_policy is missed for attribute
```

```
    MT76_TM_ATTR_TX_LENGTH. This patch adds the correct description to make
    sure the

      u32 val = nla_get_u32(tb[MT76_TM_ATTR_TX_LENGTH]);

    in function mt76_testmode_cmd() is safe and will not result in
    out-of-attribute read.

    Fixes: f0efa8621550 ("mt76: add API for testmode support")
    Signed-off-by: Lin Ma <linma@zju.edu.cn>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9ec0dec0baea36dfb5cc558f0c5d3422dd12bb20
Author: Felix Fietkau <nbd@nbd.name>
Date:   Fri Jul 14 10:57:15 2023 +0200

    wifi: mt76: mt7915: remove VHT160 capability on MT7915

    [ Upstream commit 3ec5ac12ac8a4e6b1e085374325a5fbd1b650fd5 ]

    The IEEE80211_VHT_CAP_EXT_NSS_BW value already indicates support for half-NSS
    160 MHz support, so it is wrong to also advertise full 160 MHz support.

    Fixes: c2f73eacee3b ("wifi: mt76: mt7915: add back 160MHz channel width support for
MT7915")
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0e61f73e6ebc006b90d0349d253b82f4cca72c79
Author: Felix Fietkau <nbd@nbd.name>
Date:   Thu Jun 29 22:39:30 2023 +0200

    wifi: mt76: mt7915: fix capabilities in non-AP mode

    [ Upstream commit 02a894046d5ab7d0010f39ea54fde7e167919d04 ]

    Capabilities in vif->bss_conf are only initialized in AP mode.
    For other modes, they should be enabled by default, in order to avoid a
    mismatch.

    Fixes: 885f7af7e544 ("wifi: mt76: mt7915: remove mt7915_mcu_beacon_check_caps()")
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6bce28ce2839001fd15211582bda306c0c53ceb3
Author: Rany Hany <rany_hany@riseup.net>
Date:   Sun May 28 21:04:28 2023 +0300

    wifi: mt76: mt7915: fix command timeout in AP stop period

    [ Upstream commit c4f0755823045b66484fb53d686f85d3151400f4 ]

    Due to AP stop improperly, mt7915 driver would face random command timeout
    by chip fw problem. Migrate AP start/stop process to .start_ap/.stop_ap and
    congiure BSS network settings in both hooks.

    The new flow is shown below.
    * AP start
        .start_ap()
          configure BSS network resource
          set BSS to connected state
        .bss_info_changed()
          enable fw beacon offload
```

```
      * AP stop
          .bss_info_changed()
            disable fw beacon offload (skip this command)
          .stop_ap()
            set BSS to disconnected state (beacon offload disabled automatically)
            destroy BSS network resource

      Based on "mt76: mt7921: fix command timeout in AP stop period"

      Signed-off-by: Rany Hany <rany_hany@riseup.net>
      Signed-off-by: Felix Fietkau <nbd@nbd.name>
      Stable-dep-of: 02a894046d5a ("wifi: mt76: mt7915: fix capabilities in non-AP mode")
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 3d4522f59fb748a54446846522941a4f09da63e9
  Author: Lorenz Bauer <lmb@isovalent.com>
  Date:   Thu Jul 20 17:30:06 2023 +0200

      bpf: reject unhashed sockets in bpf_sk_assign

      [ Upstream commit 67312adc96b5a585970d03b62412847afe2c6b01 ]

      The semantics for bpf_sk_assign are as follows:

          sk = some_lookup_func()
          bpf_sk_assign(skb, sk)
          bpf_sk_release(sk)

      That is, the sk is not consumed by bpf_sk_assign. The function
      therefore needs to make sure that sk lives long enough to be
      consumed from __inet_lookup_skb. The path through the stack for a
      TCPv4 packet is roughly:

        netif_receive_skb_core: takes RCU read lock
          __netif_receive_skb_core:
            sch_handle_ingress:
              tcf_classify:
                bpf_sk_assign()
            deliver_ptype_list_skb:
              deliver_skb:
                ip_packet_type->func == ip_rcv:
                  ip_rcv_core:
                  ip_rcv_finish_core:
                    dst_input:
                      ip_local_deliver:
                        ip_local_deliver_finish:
                          ip_protocol_deliver_rcu:
                            tcp_v4_rcv:
                              __inet_lookup_skb:
                                skb_steal_sock

      The existing helper takes advantage of the fact that everything
      happens in the same RCU critical section: for sockets with
      SOCK_RCU_FREE set bpf_sk_assign never takes a reference.
      skb_steal_sock then checks SOCK_RCU_FREE again and does sock_put
      if necessary.

      This approach assumes that SOCK_RCU_FREE is never set on a sk
      between bpf_sk_assign and skb_steal_sock, but this invariant is
      violated by unhashed UDP sockets. A new UDP socket is created
      in TCP_CLOSE state but without SOCK_RCU_FREE set. That flag is only
      added in udp_lib_get_port() which happens when a socket is bound.

      When bpf_sk_assign was added it wasn't possible to access unhashed
      UDP sockets from BPF, so this wasn't a problem. This changed
```

```
    in commit 0c48eefae712 ("sock_map: Lift socket state restriction
    for datagram sockets"), but the helper wasn't adjusted accordingly.
    The following sequence of events will therefore lead to a refcount
    leak:

    1. Add socket(AF_INET, SOCK_DGRAM) to a sockmap.
    2. Pull socket out of sockmap and bpf_sk_assign it. Since
       SOCK_RCU_FREE is not set we increment the refcount.
    3. bind() or connect() the socket, setting SOCK_RCU_FREE.
    4. skb_steal_sock will now set refcounted = false due to
       SOCK_RCU_FREE.
    5. tcp_v4_rcv() skips sock_put().

    Fix the problem by rejecting unhashed sockets in bpf_sk_assign().
    This matches the behaviour of __inet_lookup_skb which is ultimately
    the goal of bpf_sk_assign().

    Fixes: cf7fbe660f2d ("bpf: Add socket assign support")
    Cc: Joe Stringer <joe@cilium.io>
    Signed-off-by: Lorenz Bauer <lmb@isovalent.com>
    Reviewed-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Link: https://lore.kernel.org/r/20230720-so-reuseport-v6-2-7021b683cdae@isovalent.com
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 03760c98d5d672ce844f4dc093a697f37390477e
Author: Lorenz Bauer <lmb@isovalent.com>
Date:   Thu Jul 20 17:30:05 2023 +0200

    udp: re-score reuseport groups when connected sockets are present

    [ Upstream commit f0ea27e7bfe1c34e1f451a63eb68faa1d4c3a86d ]

    Contrary to TCP, UDP reuseport groups can contain TCP_ESTABLISHED
    sockets. To support these properly we remember whether a group has
    a connected socket and skip the fast reuseport early-return. In
    effect we continue scoring all reuseport sockets and then choose the
    one with the highest score.

    The current code fails to re-calculate the score for the result of
    lookup_reuseport. According to Kuniyuki Iwashima:

        1) SO_INCOMING_CPU is set
           -> selected sk might have +1 score

        2) BPF prog returns ESTABLISHED and/or SO_INCOMING_CPU sk
           -> selected sk will have more than 8

      Using the old score could trigger more lookups depending on the
      order that sockets are created.

        sk -> sk (SO_INCOMING_CPU) -> sk (ESTABLISHED)
        |     |
        `-> select the next SO_INCOMING_CPU sk
              |
              `-> select itself (We should save this lookup)

    Fixes: efc6b6f6c311 ("udp: Improve load balancing for SO_REUSEPORT.")
    Reviewed-by: Kuniyuki Iwashima <kuniyu@amazon.com>
    Signed-off-by: Lorenz Bauer <lmb@isovalent.com>
    Link: https://lore.kernel.org/r/20230720-so-reuseport-v6-1-7021b683cdae@isovalent.com
    Signed-off-by: Martin KaFai Lau <martin.lau@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ed05e0fbc9a7add3d49f8abc7d0ba6fe035a343f
```

```
Author: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
Date:   Mon Jun 5 23:21:39 2023 +0800

    wifi: mt76: mt7996: fix WA event ring size

    [ Upstream commit 1634de418b3048c5f435b6ffd37f75943c554c04 ]

    Fix rx ring size of WA event to get rid of event loss and queue overflow
    problems.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 08fd174ddbf7d8d4b075d646801479d34ac9c12d
Author: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
Date:   Mon Jun 5 23:21:38 2023 +0800

    wifi: mt76: mt7996: use correct phy for background radar event

    [ Upstream commit 9ffe0d5690ed916e09baad2cc9ee7ec65b110038 ]

    If driver directly uses the band_idx reported from the radar event to
    access mt76_phy array, it will get the wrong phy for background radar.
    Fix this by adjusting the statement.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Signed-off-by: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7b15d7c265f705310023f9852ef73717d31e8f26
Author: Peter Chiu <chui-hao.chiu@mediatek.com>
Date:   Mon Jun 5 23:21:34 2023 +0800

    wifi: mt76: mt7996: fix bss wlan_idx when sending bss_info command

    [ Upstream commit cc945b546227423488fe4be0ab92fd126b703246 ]

    The bmc_tx_wlan_idx should be the wlan_idx of the current bss rather
    than peer AP's wlan_idx, otherwise there will appear some frame
    decryption problems on station mode.

    Fixes: 98686cd21624 ("wifi: mt76: mt7996: add driver for MediaTek Wi-Fi 7 (802.11be)
devices")
    Reviewed-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Peter Chiu <chui-hao.chiu@mediatek.com>
    Signed-off-by: Shayne Chen <shayne.chen@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9944be4108c92ea3f4fbeef7b8c460c4cead7411
Author: Ming Yen Hsieh <mingyen.hsieh@mediatek.com>
Date:   Thu May 18 22:08:14 2023 +0800

    wifi: mt76: mt7921: fix non-PSC channel scan fail

    [ Upstream commit 0e5911bb7cc92c00dda9b4d635c1266b7ca915c6 ]

    Due to the scan command may only request legacy bands and PSC channel
```

```
    in 6GHz band, we are unable to scan the APs on non-PSC channel in this
    case. Enable WIPHY_FLAG_SPLIT_SCAN_6GHZ to support non-PSC channel
    (obtained during scan on legacy bands) in 6GHz scan request.

    Fixes: 50ac15a511e3 ("mt76: mt7921: add 6GHz support")
    Signed-off-by: Ming Yen Hsieh <mingyen.hsieh@mediatek.com>
    Signed-off-by: Deren Wu <deren.wu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 7af917d4864c6166ae7ccb9ed2e55f61867d1718
Author: Peter Chiu <chui-hao.chiu@mediatek.com>
Date:   Wed May 10 12:53:18 2023 +0800

    wifi: mt76: mt7915: rework tx bytes counting when WED is active

    [ Upstream commit f39d499345dddb8382986fd5a2a0e84a63b1a6d5 ]

    Concurrent binding/non-binding skbs could be handled anywhere which leads
    to mixed byte counting, so switch to use PPDU TxS reporting regardless Tx
    paths when WED is active.

    Fixes: 43eaa3689507 ("wifi: mt76: add PPDU based TxS support for WED device")
    Co-developed-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Peter Chiu <chui-hao.chiu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit feae00c6468ce11962012fb5f302729b76550de9
Author: Peter Chiu <chui-hao.chiu@mediatek.com>
Date:   Wed May 10 12:53:17 2023 +0800

    wifi: mt76: mt7915: rework tx packets counting when WED is active

    [ Upstream commit 161a7528e4074d104305fc109c16134b4990070e ]

    PPDU TxS can only report MPDU count whereas mac80211 requires MSDU scale
    (NL80211_STA_INFO_TX_PACKETS), so switch to get MSDU counts from WA
    statistic.

    Note that mt7915 WA firmware only counts tx_packet for WED path, so driver
    needs to take care of host path additionally.

    Fixes: 43eaa3689507 ("wifi: mt76: add PPDU based TxS support for WED device")
    Co-developed-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Peter Chiu <chui-hao.chiu@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70bbcc4ad654430537bc1530eaf0c6a9e1a75858
Author: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
Date:   Tue May 9 11:11:57 2023 +0800

    wifi: mt76: mt7915: fix background radar event being blocked

    [ Upstream commit 9a3994077d170ec9ac75e800932b5671d9940cd2 ]

    The background radar uses MT_RX_SEL2 as its band indication, so fix it.

    Fixes: 7a53eecd5c87 (wifi: mt76: mt7915: check the correctness of event data)
    Signed-off-by: StanleyYP Wang <StanleyYP.Wang@mediatek.com>
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
```

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

commit 6c5fd5f842307c60f36456ef40c45ed03bf37b48
Author: Ryder Lee <ryder.lee@mediatek.com>
Date:   Thu Apr 27 07:05:15 2023 +0800

    wifi: mt76: mt7996: fix header translation logic

    [ Upstream commit c55b4e788f1dd6ca89cc97cf291d2a03b0b96de1 ]

    When header translation failure is indicated, the hardware will insert
    an extra 2-byte field containing the data length after the protocol
    type field. This happens either when the LLC-SNAP pattern did not match,
    or if a VLAN header was detected.

    The previous commit accidentally breaks the logic, so reverts back.

    Fixes: 27db47ab1f47 (wifi: mt76: mt7996: enable mesh HW amsdu/de-amsdu support)
    Signed-off-by: Ryder Lee <ryder.lee@mediatek.com>
    Signed-off-by: Felix Fietkau <nbd@nbd.name>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9a3e0d7248b2e6bbae996f599e5ebd15bdf51181
Author: Zhang Shurong <zhang_shurong@foxmail.com>
Date:   Sat Jul 15 21:42:57 2023 +0800

    wifi: rtw89: debug: Fix error handling in rtw89_debug_priv_btc_manual_set()

    [ Upstream commit 59b4cc439f184c5eaa34161ec67af1e16ffabed4 ]

    If there is a failure during kstrtobool_from_user()
    rtw89_debug_priv_btc_manual_set should return a negative error code
    instead of returning the count directly.

    Fix this bug by returning an error code instead of a count after
    a failed call of the function "kstrtobool_from_user". Moreover
    I omitted the label "out" with this source code correction.

    Fixes: e3ec7017f6a2 ("rtw89: add Realtek 802.11ax driver")
    Signed-off-by: Zhang Shurong <zhang_shurong@foxmail.com>
    Acked-by: Ping-Ke Shih <pkshih@realtek.com>
    Signed-off-by: Kalle Valo <kvalo@kernel.org>
    Link: https://lore.kernel.org/r/tencent_1C09B99BD7DA9CAD18B00C8F0F050F540607@qq.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bcec8481ee555817620faac835954156eca08695
Author: Dan Carpenter <dan.carpenter@linaro.org>
Date:   Fri Jul 21 17:55:33 2023 +0300

    regmap: rbtree: Use alloc_flags for memory allocations

    [ Upstream commit 0c8b0bf42c8cef56f7cd9cd876fbb7ece9217064 ]

    The kunit tests discovered a sleeping in atomic bug.  The allocations
    in the regcache-rbtree code should use the map->alloc_flags instead of
    GFP_KERNEL.

    [    5.005510] BUG: sleeping function called from invalid context at
include/linux/sched/mm.h:306
    [    5.005960] in_atomic(): 1, irqs_disabled(): 128, non_block: 0, pid: 117, name:
kunit_try_catch
    [    5.006219] preempt_count: 1, expected: 0
    [    5.006414] 1 lock held by kunit_try_catch/117:
    [    5.006590]  #0: 833b9010 (regmap_kunit:86:(config)->lock){....}-{2:2}, at:
regmap_lock_spinlock+0x14/0x1c
```

```
     [    5.007493] irq event stamp: 162
     [    5.007627] hardirqs last  enabled at (161): [<80786738>] crng_make_state+0x1a0/0x294
     [    5.007871] hardirqs last disabled at (162): [<80c531ec>]
 _raw_spin_lock_irqsave+0x7c/0x80
     [    5.008119] softirqs last  enabled at (0): [<801110ac>] copy_process+0x810/0x2138
     [    5.008356] softirqs last disabled at (0): [<00000000>] 0x0
     [    5.008688] CPU: 0 PID: 117 Comm: kunit_try_catch Tainted: G                N 6.4.4-
 rc3-g0e8d2fdfb188 #1
     [    5.009011] Hardware name: Generic DT based system
     [    5.009277]  unwind_backtrace from show_stack+0x18/0x1c
     [    5.009497]  show_stack from dump_stack_lvl+0x38/0x5c
     [    5.009676]  dump_stack_lvl from __might_resched+0x188/0x2d0
     [    5.009860]  __might_resched from __kmem_cache_alloc_node+0x1dc/0x25c
     [    5.010061]  __kmem_cache_alloc_node from kmalloc_trace+0x30/0xc8
     [    5.010254]  kmalloc_trace from regcache_rbtree_write+0x26c/0x468
     [    5.010446]  regcache_rbtree_write from _regmap_write+0x88/0x140
     [    5.010634]  _regmap_write from regmap_write+0x44/0x68
     [    5.010803]  regmap_write from basic_read_write+0x8c/0x270
     [    5.010980]  basic_read_write from kunit_try_run_case+0x48/0xa0

     Fixes: 28644c809f44 ("regmap: Add the rbtree cache support")
     Reported-by: Guenter Roeck <linux@roeck-us.net>
     Closes: https://lore.kernel.org/all/ee59d128-413c-48ad-a3aa-d9d350c80042@roeck-us.net/
     Signed-off-by: Dan Carpenter <dan.carpenter@linaro.org>
     Tested-by: Guenter Roeck <linux@roeck-us.net>
     Link: https://lore.kernel.org/r/58f12a07-5f4b-4a8f-ab84-0a42d1908cb9@moroto.mountain
     Signed-off-by: Mark Brown <broonie@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 4e0ad6282cc37847674b5258240d4ede148058f8
 Author: Guenter Roeck <linux@roeck-us.net>
 Date:   Thu Jul 20 10:20:21 2023 -0700

     regmap: maple: Use alloc_flags for memory allocations

     [ Upstream commit b0393e1fe40e962574613a5cdc4a470d6c1de023 ]

     REGCACHE_MAPLE needs to allocate memory for regmap operations.
     This results in lockdep splats if used with fast_io since fast_io uses
     spinlocks for locking.

     BUG: sleeping function called from invalid context at include/linux/sched/mm.h:306
     in_atomic(): 1, irqs_disabled(): 128, non_block: 0, pid: 167, name: kunit_try_catch
     preempt_count: 1, expected: 0
     1 lock held by kunit_try_catch/167:
      #0: 838e9c10 (regmap_kunit:86:(config)->lock){....}-{2:2}, at:
 regmap_lock_spinlock+0x14/0x1c
     irq event stamp: 146
     hardirqs last  enabled at (145): [<8078bfa8>] crng_make_state+0x1a0/0x294
     hardirqs last disabled at (146): [<80c5f62c>] _raw_spin_lock_irqsave+0x7c/0x80
     softirqs last  enabled at (0): [<80110cc4>] copy_process+0x810/0x216c
     softirqs last disabled at (0): [<00000000>] 0x0
     CPU: 0 PID: 167 Comm: kunit_try_catch Tainted: G                N 6.5.0-rc1-00028-
 gc4be22597a36-dirty #6
     Hardware name: Generic DT based system
      unwind_backtrace from show_stack+0x18/0x1c
      show_stack from dump_stack_lvl+0x38/0x5c
      dump_stack_lvl from __might_resched+0x188/0x2d0
      __might_resched from __kmem_cache_alloc_node+0x1f4/0x258
      __kmem_cache_alloc_node from __kmalloc+0x48/0x170
      __kmalloc from regcache_maple_write+0x194/0x248
      regcache_maple_write from _regmap_write+0x88/0x140
      _regmap_write from regmap_write+0x44/0x68
      regmap_write from basic_read_write+0x8c/0x27c
      basic_read_write from kunit_generic_run_threadfn_adapter+0x1c/0x28
```

```
      kunit_generic_run_threadfn_adapter from kthread+0xf8/0x120
      kthread from ret_from_fork+0x14/0x3c
     Exception stack(0x881a5fb0 to 0x881a5ff8)
     5fa0:                                     00000000 00000000 00000000 00000000
     5fc0: 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
     5fe0: 00000000 00000000 00000000 00000000 00000013 00000000

     Use map->alloc_flags instead of GFP_KERNEL for memory allocations to fix
     the problem.

     Fixes: f033c26de5a5 ("regmap: Add maple tree based register cache")
     Cc: Dan Carpenter <dan.carpenter@linaro.org>
     Signed-off-by: Guenter Roeck <linux@roeck-us.net>
     Link: https://lore.kernel.org/r/20230720172021.2617326-1-linux@roeck-us.net
     Signed-off-by: Mark Brown <broonie@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 69dfa8c7e88a2997bb2e3dd8d91d03f51c5a175e
Author: Martin Kaiser <martin@kaiser.cx>
Date:   Tue Jul 4 19:32:01 2023 +0200

    hwrng: pic32 - use devm_clk_get_enabled

    [ Upstream commit 6755ad74aac0fb1c79b14724feb81b2f6ff25847 ]

    Use devm_clk_get_enabled in the pic32 driver. Ensure that the clock is
    enabled as long as the driver is registered with the hwrng core.

    Fixes: 7ea39973d1e5 ("hwrng: pic32 - Use device-managed registration API")
    Signed-off-by: Martin Kaiser <martin@kaiser.cx>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9681618bf379566285fe451735ec20a1b19faffb
Author: Martin Kaiser <martin@kaiser.cx>
Date:   Sun Jul 2 19:35:02 2023 +0200

    hwrng: nomadik - keep clock enabled while hwrng is registered

    [ Upstream commit 039980de89dc9dd757418d6f296e4126cc3f86c3 ]

    The nomadik driver uses devres to register itself with the hwrng core,
    the driver will be unregistered from hwrng when its device goes out of
    scope. This happens after the driver's remove function is called.

    However, nomadik's clock is disabled in the remove function. There's a
    short timeframe where nomadik is still registered with the hwrng core
    although its clock is disabled. I suppose the clock must be active to
    access the hardware and serve requests from the hwrng core.

    Switch to devm_clk_get_enabled and let devres disable the clock and
    unregister the hwrng. This avoids the race condition.

    Fixes: 3e75241be808 ("hwrng: drivers - Use device-managed registration API")
    Signed-off-by: Martin Kaiser <martin@kaiser.cx>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a54f60bbbf84efae424cfc7e99efbac1e0fbf7ad
Author: Eric Dumazet <edumazet@google.com>
Date:   Tue Jul 18 16:20:49 2023 +0000

    tcp: tcp_enter_quickack_mode() should be static

    [ Upstream commit 03b123debcbc8db987bda17ed8412cc011064c22 ]
```

```
    After commit d2ccd7bc8acd ("tcp: avoid resetting ACK timer in DCTCP"),
    tcp_enter_quickack_mode() is only used from net/ipv4/tcp_input.c.

    Fixes: d2ccd7bc8acd ("tcp: avoid resetting ACK timer in DCTCP")
    Signed-off-by: Eric Dumazet <edumazet@google.com>
    Cc: Yuchung Cheng <ycheng@google.com>
    Cc: Neal Cardwell <ncardwell@google.com>
    Link: https://lore.kernel.org/r/20230718162049.1444938-1-edumazet@google.com
    Signed-off-by: Jakub Kicinski <kuba@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aa39f43a498e766aff12b8b39eeff8f1450d2c8a
Author: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
Date:   Thu Jun 22 10:26:35 2023 +0100

    crypto: qat - change value of default idle filter

    [ Upstream commit 0f942bdfe9d463be3073301519492f8d53c6b2d5 ]

    The power management configuration of 4xxx devices is too aggressive
    and in some conditions the device might be prematurely put to a low
    power state.
    Increase the idle filter value to prevent that.
    In future, this will be set by firmware.

    Fixes: e5745f34113b ("crypto: qat - enable power management for QAT GEN4")
    Signed-off-by: Giovanni Cabiddu <giovanni.cabiddu@intel.com>
    Reviewed-by: Damian Muszynski <damian.muszynski@intel.com>
    Signed-off-by: Herbert Xu <herbert@gondor.apana.org.au>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3b3a08048592d436deac7fe9048610be224435e
Author: Yafang Shao <laoar.shao@gmail.com>
Date:   Thu Jul 13 02:56:41 2023 +0000

    bpf: Fix an error in verifying a field in a union

    [ Upstream commit 33937607efa050d9e237e0c4ac4ada02d961c466 ]

    We are utilizing BPF LSM to monitor BPF operations within our container
    environment. When we add support for raw_tracepoint, it hits below
    error.

    ; (const void *)attr->raw_tracepoint.name);
    27: (79) r3 = *(u64 *)(r2 +0)
    access beyond the end of member map_type (mend:4) in struct (anon) with off 0 size 8

    It can be reproduced with below BPF prog.

    SEC("lsm/bpf")
    int BPF_PROG(bpf_audit, int cmd, union bpf_attr *attr, unsigned int size)
    {
            switch (cmd) {
            case BPF_RAW_TRACEPOINT_OPEN:
                    bpf_printk("raw_tracepoint is %s", attr->raw_tracepoint.name);
                    break;
            default:
                    break;
            }
            return 0;
    }

    The reason is that when accessing a field in a union, such as bpf_attr,
    if the field is located within a nested struct that is not the first
```

```
    member of the union, it can result in incorrect field verification.

      union bpf_attr {
          struct {
              __u32 map_type; <<<< Actually it will find that field.
              __u32 key_size;
              __u32 value_size;
             ...
          };
          ...
          struct {
              __u64 name;    <<<< We want to verify this field.
              __u32 prog_fd;
          } raw_tracepoint;
      };

    Considering the potential deep nesting levels, finding a perfect
    solution to address this issue has proven challenging. Therefore, I
    propose a solution where we simply skip the verification process if the
    field in question is located within a union.

    Fixes: 7e3617a72df3 ("bpf: Add array support to btf_struct_access")
    Signed-off-by: Yafang Shao <laoar.shao@gmail.com>
    Link: https://lore.kernel.org/r/20230713025642.27477-4-laoar.shao@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 208744e8cb224bf1d23bf22da323cc2fa1ef8c76
Author: Yafang Shao <laoar.shao@gmail.com>
Date:   Thu Jul 13 02:56:39 2023 +0000

    bpf: Fix an error around PTR_UNTRUSTED

    [ Upstream commit 7ce4dc3e4a9d954c8a1fb483c7a527e9b060b860 ]

    Per discussion with Alexei, the PTR_UNTRUSTED flag should not been
    cleared when we start to walk a new struct, because the struct in
    question may be a struct nested in a union. We should also check and set
    this flag before we walk its each member, in case itself is a union.
    We will clear this flag if the field is BTF_TYPE_SAFE_RCU_OR_NULL.

    Fixes: 6fcd486b3a0a ("bpf: Refactor RCU enforcement in the verifier.")
    Signed-off-by: Yafang Shao <laoar.shao@gmail.com>
    Link: https://lore.kernel.org/r/20230713025642.27477-2-laoar.shao@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit ee55daf6f28519f188fbc719a1e76200239726ad
Author: Yafang Shao <laoar.shao@gmail.com>
Date:   Sun Jul 9 02:56:25 2023 +0000

    bpf: Clear the probe_addr for uprobe

    [ Upstream commit 5125e757e62f6c1d5478db4c2b61a744060ddf3f ]

    To avoid returning uninitialized or random values when querying the file
    descriptor (fd) and accessing probe_addr, it is necessary to clear the
    variable prior to its use.

    Fixes: 41bdc4b40ed6 ("bpf: introduce bpf subcommand BPF_TASK_FD_QUERY")
    Signed-off-by: Yafang Shao <laoar.shao@gmail.com>
    Acked-by: Yonghong Song <yhs@fb.com>
    Acked-by: Jiri Olsa <jolsa@kernel.org>
    Link: https://lore.kernel.org/r/20230709025630.3735-6-laoar.shao@gmail.com
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
```

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

commit 2ab7b3575154ffcf66ce0aaafc1d5d058252f40b
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Mon Jul 10 19:41:50 2023 -0700

    libbpf: Fix realloc API handling in zero-sized edge cases

    [ Upstream commit 8a0260dbf6553c969248b6530cafadac46562f47 ]

    realloc() and reallocarray() can either return NULL or a special
    non-NULL pointer, if their size argument is zero. This requires a bit
    more care to handle NULL-as-valid-result situation differently from
    NULL-as-error case. This has caused real issues before ([0]), and just
    recently bit again in production when performing bpf_program__attach_usdt().

    This patch fixes 4 places that do or potentially could suffer from this
    mishandling of NULL, including the reported USDT-related one.

    There are many other places where realloc()/reallocarray() is used and
    NULL is always treated as an error value, but all those have guarantees
    that their size is always non-zero, so those spot don't need any extra
    handling.

      [0] d08ab82f59d5 ("libbpf: Fix double-free when linker processes empty sections")

    Fixes: 999783c8bbda ("libbpf: Wire up spec management and other arch-independent USDT
logic")
    Fixes: b63b3c490eee ("libbpf: Add bpf_program__set_insns function")
    Fixes: 697f104db8a6 ("libbpf: Support custom SEC() handlers")
    Fixes: b12688267280 ("libbpf: Change the order of data and text relocations.")
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/20230711024150.1566433-1-andrii@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e03118c52f971d5641dcb1941c75ddab59acfac8
Author: Alexander Lobakin <alobakin@pm.me>
Date:   Fri Jul 7 10:54:25 2023 +0100

    bpftool: Use a local bpf_perf_event_value to fix accessing its fields

    [ Upstream commit 658ac06801315b739774a15796ff06913ef5cad5 ]

    Fix the following error when building bpftool:

      CLANG   profiler.bpf.o
      CLANG   pid_iter.bpf.o
    skeleton/profiler.bpf.c:18:21: error: invalid application of 'sizeof' to an incomplete
type 'struct bpf_perf_event_value'
           __uint(value_size, sizeof(struct bpf_perf_event_value));
                              ^     ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    tools/bpf/bpftool/bootstrap/libbpf/include/bpf/bpf_helpers.h:13:39: note: expanded from
macro '__uint'
    tools/bpf/bpftool/bootstrap/libbpf/include/bpf/bpf_helper_defs.h:7:8: note: forward
declaration of 'struct bpf_perf_event_value'
    struct bpf_perf_event_value;
           ^

    struct bpf_perf_event_value is being used in the kernel only when
    CONFIG_BPF_EVENTS is enabled, so it misses a BTF entry then.
    Define struct bpf_perf_event_value___local with the
    `preserve_access_index` attribute inside the pid_iter BPF prog to
    allow compiling on any configs. It is a full mirror of a UAPI
    structure, so is compatible both with and w/o CO-RE.
```

```
    bpf_perf_event_read_value() requires a pointer of the original type,
    so a cast is needed.

    Fixes: 47c09d6a9f67 ("bpftool: Introduce "prog profile" command")
    Suggested-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexander Lobakin <alobakin@pm.me>
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20230707095425.168126-5-quentin@isovalent.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c6e7ee088b223bccd047db802126faf1cd5eef1c
Author: Quentin Monnet <quentin@isovalent.com>
Date:   Fri Jul 7 10:54:24 2023 +0100

    bpftool: Use a local copy of BPF_LINK_TYPE_PERF_EVENT in pid_iter.bpf.c

    [ Upstream commit 44ba7b30e84fb40da2295e85a6d209e199fdc977 ]

    In order to allow the BPF program in bpftool's pid_iter.bpf.c to compile
    correctly on hosts where vmlinux.h does not define
    BPF_LINK_TYPE_PERF_EVENT (running kernel versions lower than 5.15, for
    example), define and use a local copy of the enum value. This requires
    LLVM 12 or newer to build the BPF program.

    Fixes: cbdaf71f7e65 ("bpftool: Add bpf_cookie to link output")
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20230707095425.168126-4-quentin@isovalent.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bb4c88ff2570ce90543048cdd6da7e3ddc013914
Author: Alexander Lobakin <alobakin@pm.me>
Date:   Fri Jul 7 10:54:23 2023 +0100

    bpftool: Define a local bpf_perf_link to fix accessing its fields

    [ Upstream commit 67a43462ee2405c94e985a747bdcb8e3a0d66203 ]

    When building bpftool with !CONFIG_PERF_EVENTS:

    skeleton/pid_iter.bpf.c:47:14: error: incomplete definition of type 'struct
bpf_perf_link'
            perf_link = container_of(link, struct bpf_perf_link, link);
                        ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
    tools/bpf/bpftool/bootstrap/libbpf/include/bpf/bpf_helpers.h:74:22: note: expanded from
macro 'container_of'
                    ((type *)(__mptr - offsetof(type, member)));    \
                                       ^~~~~~~~~~~~~~~~~~~~~~
    tools/bpf/bpftool/bootstrap/libbpf/include/bpf/bpf_helpers.h:68:60: note: expanded from
macro 'offsetof'
     #define offsetof(TYPE, MEMBER)  ((unsigned long)&((TYPE *)0)->MEMBER)
                                                      ~~~~~~~~~~~^
    skeleton/pid_iter.bpf.c:44:9: note: forward declaration of 'struct bpf_perf_link'
            struct bpf_perf_link *perf_link;
                   ^

    &bpf_perf_link is being defined and used only under the ifdef.
    Define struct bpf_perf_link___local with the `preserve_access_index`
    attribute inside the pid_iter BPF prog to allow compiling on any
    configs. CO-RE will substitute it with the real struct bpf_perf_link
    accesses later on.
    container_of() uses offsetof(), which does the necessary CO-RE
    relocation if the field is specified with `preserve_access_index` - as
    is the case for struct bpf_perf_link___local.
```

```
    Fixes: cbdaf71f7e65 ("bpftool: Add bpf_cookie to link output")
    Suggested-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexander Lobakin <alobakin@pm.me>
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20230707095425.168126-3-quentin@isovalent.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 3abfaffa166e1910eccdebdd110a9c0fa9de9537
Author: Alexander Lobakin <alobakin@pm.me>
Date:   Fri Jul 7 10:54:22 2023 +0100

    bpftool: use a local copy of perf_event to fix accessing :: Bpf_cookie

    [ Upstream commit 4cbeeb0dc02f8ac7b975b2ab0080ace53d43d62a ]

    When CONFIG_PERF_EVENTS is not set, struct perf_event remains empty.
    However, the structure is being used by bpftool indirectly via BTF.
    This leads to:

    skeleton/pid_iter.bpf.c:49:30: error: no member named 'bpf_cookie' in 'struct
perf_event'
            return BPF_CORE_READ(event, bpf_cookie);
                   ~~~~~~~~~~~~~~~~~~~~~^~~~~~~~~~~


    ...

    skeleton/pid_iter.bpf.c:49:9: error: returning 'void' from a function with incompatible
result type '__u64' (aka 'unsigned long long')
            return BPF_CORE_READ(event, bpf_cookie);
                   ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

    Tools and samples can't use any CONFIG_ definitions, so the fields
    used there should always be present.
    Define struct perf_event___local with the `preserve_access_index`
    attribute inside the pid_iter BPF prog to allow compiling on any
    configs. CO-RE will substitute it with the real struct perf_event
    accesses later on.

    Fixes: cbdaf71f7e65 ("bpftool: Add bpf_cookie to link output")
    Suggested-by: Andrii Nakryiko <andrii@kernel.org>
    Signed-off-by: Alexander Lobakin <alobakin@pm.me>
    Signed-off-by: Quentin Monnet <quentin@isovalent.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Link: https://lore.kernel.org/bpf/20230707095425.168126-2-quentin@isovalent.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f84c4ac69873b86e7532f922749e5fd2baace6a
Author: Andrii Nakryiko <andrii@kernel.org>
Date:   Fri Jul 7 16:11:56 2023 -0700

    libbpf: only reset sec_def handler when necessary

    [ Upstream commit c628747cc8800cf6d33d09f7f42c8b6f91e64dc7 ]

    Don't reset recorded sec_def handler unconditionally on
    bpf_program__set_type(). There are two situations where this is wrong.

    First, if the program type didn't actually change. In that case original
    SEC handler should work just fine.

    Second, catch-all custom SEC handler is supposed to work with any BPF
    program type and SEC() annotation, so it also doesn't make sense to
    reset that.
```

```
    This patch fixes both issues. This was reported recently in the context
    of breaking perf tool, which uses custom catch-all handler for fancy BPF
    prologue generation logic. This patch should fix the issue.

      [0] https://lore.kernel.org/linux-perf-users/ab865e6d-06c5-078e-e404-
7f90686db50d@amd.com/

    Fixes: d6e6286a12e7 ("libbpf: disassociate section handler on explicit
bpf_program__set_type() call")
    Reported-by: Ravi Bangoria <ravi.bangoria@amd.com>
    Signed-off-by: Andrii Nakryiko <andrii@kernel.org>
    Acked-by: Stanislav Fomichev <sdf@google.com>
    Link: https://lore.kernel.org/r/20230707231156.1711948-1-andrii@kernel.org
    Signed-off-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f618dee0f27bb57b54ed6bcbee6752d953a502b1
Author: Daniel Borkmann <daniel@iogearbox.net>
Date:   Mon Jun 26 15:19:42 2023 +0200

    selftests/bpf: Fix bpf_nf failure upon test rerun

    [ Upstream commit 17e8e5d6e09adb4b4f4fb5c89b3ec3fcae2c64a6 ]

    Alexei reported:

      After fast forwarding bpf-next today bpf_nf test started to fail when
      run twice:

      $ ./test_progs -t bpf_nf
      #17      bpf_nf:OK
      Summary: 1/10 PASSED, 0 SKIPPED, 0 FAILED

      $ ./test_progs -t bpf_nf
      All error logs:
      test_bpf_nf_ct:PASS:test_bpf_nf__open_and_load 0 nsec
      test_bpf_nf_ct:PASS:iptables-legacy -t raw -A PREROUTING -j CONNMARK
      --set-mark 42/0 0 nsec
      (network_helpers.c:102: errno: Address already in use) Failed to bind socket
      test_bpf_nf_ct:FAIL:start_server unexpected start_server: actual -1 < expected 0
      #17/1    bpf_nf/xdp-ct:FAIL
      test_bpf_nf_ct:PASS:test_bpf_nf__open_and_load 0 nsec
      test_bpf_nf_ct:PASS:iptables-legacy -t raw -A PREROUTING -j CONNMARK
      --set-mark 42/0 0 nsec
      (network_helpers.c:102: errno: Address already in use) Failed to bind socket
      test_bpf_nf_ct:FAIL:start_server unexpected start_server: actual -1 < expected 0
      #17/2    bpf_nf/tc-bpf-ct:FAIL
      #17      bpf_nf:FAIL
      Summary: 0/8 PASSED, 0 SKIPPED, 1 FAILED

    I was able to locally reproduce as well. Rearrange the connection teardown
    so that the client closes its connection first so that we don't need to
    linger in TCP time-wait.

    Fixes: e81fbd4c1ba7 ("selftests/bpf: Add existing connection bpf_*_ct_lookup() test")
    Reported-by: Alexei Starovoitov <ast@kernel.org>
    Signed-off-by: Daniel Borkmann <daniel@iogearbox.net>
    Link: https://lore.kernel.org/bpf/CAADnVQ+0dnDq_v_vH1EfkacbfGnHANaon7zsw10pMb-
D9FS0Pw@mail.gmail.com
    Link: https://lore.kernel.org/bpf/20230626131942.5100-1-daniel@iogearbox.net
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 002a658fa235ee2c1af7f6a0a2162c2a6dff9027
Author: Sumit Gupta <sumitg@nvidia.com>
```

```
Date:   Fri Aug 25 16:46:17 2023 +0530

    cpufreq: tegra194: remove opp table in exit hook

    [ Upstream commit de0e85b29edfc68046d587c7d67bbd2bdc31b73f ]

    Add exit hook and remove OPP table when the device gets unregistered.
    This will fix the error messages when the CPU FREQ driver module is
    removed and then re-inserted. It also fixes these messages while
    onlining the first CPU from a policy whose all CPU's were previously
    offlined.

     debugfs: File 'cpu5' in directory 'opp' already present!
     debugfs: File 'cpu6' in directory 'opp' already present!
     debugfs: File 'cpu7' in directory 'opp' already present!

    Fixes: f41e1442ac5b ("cpufreq: tegra194: add OPP support and set bandwidth")
    Signed-off-by: Sumit Gupta <sumitg@nvidia.com>
    [ Viresh: Dropped irrelevant change from it ]
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88dbf430b635f9e5a3bf15814067e26e4a43c8b9
Author: Sumit Gupta <sumitg@nvidia.com>
Date:   Fri Aug 25 16:49:20 2023 +0530

    cpufreq: tegra194: add online/offline hooks

    [ Upstream commit a3aa97be69a7cc14ddc2bb0add0b9c51cb74bf83 ]

    Implement the light-weight tear down and bring up helpers to reduce the
    amount of work to do on CPU offline/online operation.
    This change helps to make the hotplugging paths much faster.

    Suggested-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sumit Gupta <sumitg@nvidia.com>
    Link: https://lore.kernel.org/lkml/20230816033402.3abmugb5goypvllm@vireshk-i7/
    [ Viresh: Fixed rebase conflict ]
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Stable-dep-of: de0e85b29edf ("cpufreq: tegra194: remove opp table in exit hook")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit bf6b336cf270e912f7cb4aff6b64dc0349f9b009
Author: Liao Chang <liaochang1@huawei.com>
Date:   Sat Aug 26 09:51:13 2023 +0000

    cpufreq: powernow-k8: Use related_cpus instead of cpus in driver.exit()

    [ Upstream commit 03997da042dac73c69e60d91942c727c76828b65 ]

    Since the 'cpus' field of policy structure will become empty in the
    cpufreq core API, it is better to use 'related_cpus' in the exit()
    callback of driver.

    Fixes: c3274763bfc3 ("cpufreq: powernow-k8: Initialize per-cpu data-structures
properly")
    Signed-off-by: Liao Chang <liaochang1@huawei.com>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 933e46df5bda478b64c314275dafffbb00884dfb
Author: Mikel Rychliski <mikel@mikelr.com>
Date:   Wed Aug 23 17:51:58 2023 -0400

    x86/efistub: Fix PCI ROM preservation in mixed mode
```

```
    [ Upstream commit 8b94da92559f7e403dc7ab81937cc50f949ee2fd ]

    preserve_pci_rom_image() was accessing the romsize field in
    efi_pci_io_protocol_t directly instead of using the efi_table_attr()
    helper. This prevents the ROM image from being saved correctly during a
    mixed mode boot.

    Fixes: 2c3625cb9fa2 ("efi/x86: Fold __setup_efi_pci32() and __setup_efi_pci64() into one
function")
    Signed-off-by: Mikel Rychliski <mikel@mikelr.com>
    Signed-off-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 65f40c946f7a34f25638e38b80d5ee64cde45ae1
Author: Zhangjin Wu <falcon@tinylab.org>
Date:   Sun Jul 16 02:18:54 2023 +0800

    tools/nolibc: fix up startup failures for -O0 under gcc < 11.1.0

    [ Upstream commit bff60150f7c464d80d86f289c056c2ad2afb3c05 ]

    As gcc doc [1] shows:

      Most optimizations are completely disabled at -O0 or if an -O level is
      not set on the command line, even if individual optimization flags are
      specified.

    Test result [2] shows, gcc>=11.1.0 deviates from the above description,
    but before gcc 11.1.0, "-O0" still forcely uses frame pointer in the
    _start function even if the individual optimize("omit-frame-pointer")
    flag is specified.

    The frame pointer related operations will change the stack pointer (e.g.
    In x86_64, an extra "push %rbp" will be inserted at the beginning of
    _start) and make it differs from the one we expected, as a result, break
    the whole startup function.

    To fix up this issue, as suggested by Thomas, the individual "Os" and
    "omit-frame-pointer" optimize flags are used together on _start function
    to disable frame pointer completely even if the -O0 is set on the
    command line.

    [1]: https://gcc.gnu.org/onlinedocs/gcc/Optimize-Options.html
    [2]: https://lore.kernel.org/lkml/20230714094723.140603-1-falcon@tinylab.org/

    Suggested-by: Thomas Weißschuh <linux@weissschuh.net>
    Link: https://lore.kernel.org/lkml/34b21ba5-7b59-4b3b-9ed6-ef9a3a5e06f7@t-8ch.de/
    Fixes: 7f8548589661 ("tools/nolibc: make compiler and assembler agree on the section
around _start")
    Signed-off-by: Zhangjin Wu <falcon@tinylab.org>
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9652b614d1e4e80f50d35d2af03d28e68bcb8257
Author: Zhangjin Wu <falcon@tinylab.org>
Date:   Sun Jul 16 02:17:43 2023 +0800

    tools/nolibc: arch-*.h: add missing space after ','

    [ Upstream commit 20233498359a29f7b2ff4e8fbdb0a1a7c8d5744c ]

    Fix up such errors reported by scripts/checkpatch.pl:

        ERROR: space required after that ',' (ctx:VxV)
```

```
        #148: FILE: tools/include/nolibc/arch-aarch64.h:148:
        +void __attribute__((weak,noreturn,optimize("omit-frame-pointer")))
__no_stack_protector _start(void)
                            ^

        ERROR: space required after that ',' (ctx:VxV)
        #148: FILE: tools/include/nolibc/arch-aarch64.h:148:
        +void __attribute__((weak,noreturn,optimize("omit-frame-pointer")))
__no_stack_protector _start(void)
                                  ^
```

```
    Signed-off-by: Zhangjin Wu <falcon@tinylab.org>
    Signed-off-by: Willy Tarreau <w@1wt.eu>
    Stable-dep-of: bff60150f7c4 ("tools/nolibc: fix up startup failures for -O0 under gcc <
11.1.0")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0f74f12ee042fd72e45f0e8700e063c84ef3883b
Author: Swapnil Sapkal <swapnil.sapkal@amd.com>
Date:   Fri Aug 18 11:44:52 2023 +0000

    cpufreq: amd-pstate-ut: Fix kernel panic when loading the driver

    [ Upstream commit 60dd283804479c4a52f995b713f448e2cd65b8c8 ]

    After loading the amd-pstate-ut driver, amd_pstate_ut_check_perf()
    and amd_pstate_ut_check_freq() use cpufreq_cpu_get() to get the policy
    of the CPU and mark it as busy.

    In these functions, cpufreq_cpu_put() should be used to release the
    policy, but it is not, so any other entity trying to access the policy
    is blocked indefinitely.

    One such scenario is when amd_pstate mode is changed, leading to the
    following splat:

    [ 1332.103727] INFO: task bash:2929 blocked for more than 120 seconds.
    [ 1332.110001]       Not tainted 6.5.0-rc2-amd-pstate-ut #5
    [ 1332.115315] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message.
    [ 1332.123140] task:bash            state:D stack:0     pid:2929  ppid:2873
flags:0x00004006
    [ 1332.123143] Call Trace:
    [ 1332.123145]  <TASK>
    [ 1332.123148]  __schedule+0x3c1/0x16a0
    [ 1332.123154]  ? _raw_read_lock_irqsave+0x2d/0x70
    [ 1332.123157]  schedule+0x6f/0x110
    [ 1332.123160]  schedule_timeout+0x14f/0x160
    [ 1332.123162]  ? preempt_count_add+0x86/0xd0
    [ 1332.123165]  __wait_for_common+0x92/0x190
    [ 1332.123168]  ? __pfx_schedule_timeout+0x10/0x10
    [ 1332.123170]  wait_for_completion+0x28/0x30
    [ 1332.123173]  cpufreq_policy_put_kobj+0x4d/0x90
    [ 1332.123177]  cpufreq_policy_free+0x157/0x1d0
    [ 1332.123178]  ? preempt_count_add+0x58/0xd0
    [ 1332.123180]  cpufreq_remove_dev+0xb6/0x100
    [ 1332.123182]  subsys_interface_unregister+0x114/0x120
    [ 1332.123185]  ? preempt_count_add+0x58/0xd0
    [ 1332.123187]  ? __pfx_amd_pstate_change_driver_mode+0x10/0x10
    [ 1332.123190]  cpufreq_unregister_driver+0x3b/0xd0
    [ 1332.123192]  amd_pstate_change_driver_mode+0x1e/0x50
    [ 1332.123194]  store_status+0xe9/0x180
    [ 1332.123197]  dev_attr_store+0x1b/0x30
    [ 1332.123199]  sysfs_kf_write+0x42/0x50
    [ 1332.123202]  kernfs_fop_write_iter+0x143/0x1d0
    [ 1332.123204]  vfs_write+0x2df/0x400
```

```
     [ 1332.123208]  ksys_write+0x6b/0xf0
     [ 1332.123210]  __x64_sys_write+0x1d/0x30
     [ 1332.123213]  do_syscall_64+0x60/0x90
     [ 1332.123216]  ? fpregs_assert_state_consistent+0x2e/0x50
     [ 1332.123219]  ? exit_to_user_mode_prepare+0x49/0x1a0
     [ 1332.123223]  ? irqentry_exit_to_user_mode+0xd/0x20
     [ 1332.123225]  ? irqentry_exit+0x3f/0x50
     [ 1332.123226]  ? exc_page_fault+0x8e/0x190
     [ 1332.123228]  entry_SYSCALL_64_after_hwframe+0x6e/0xd8
     [ 1332.123232] RIP: 0033:0x7fa74c514a37
     [ 1332.123234] RSP: 002b:00007ffe31dd0788 EFLAGS: 00000246 ORIG_RAX: 0000000000000001
     [ 1332.123238] RAX: ffffffffffffffda RBX: 0000000000000008 RCX: 00007fa74c514a37
     [ 1332.123239] RDX: 0000000000000008 RSI: 000055e27c447aa0 RDI: 0000000000000001
     [ 1332.123241] RBP: 000055e27c447aa0 R08: 00007fa74c5d1460 R09: 000000007fffffff
     [ 1332.123242] R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000008
     [ 1332.123244] R13: 00007fa74c61a780 R14: 00007fa74c616600 R15: 00007fa74c615a00
     [ 1332.123247]  </TASK>

     Fix this by calling cpufreq_cpu_put() wherever necessary.

     Fixes: 14eb1c96e3a3 ("cpufreq: amd-pstate: Add test module for amd-pstate driver")
     Reviewed-by: Mario Limonciello <mario.limonciello@amd.com>
     Reviewed-by: Meng Li <li.meng@amd.com>
     Reviewed-by: Wyes Karny <wyes.karny@amd.com>
     Suggested-by: Wyes Karny <wyes.karny@amd.com>
     Signed-off-by: Swapnil Sapkal <swapnil.sapkal@amd.com>
     [ rjw: Subject and changelog edits ]
     Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1b9a602bc61bcd5f25029963c5042d39990d8717
Author: Swapnil Sapkal <swapnil.sapkal@amd.com>
Date:   Fri Aug 18 11:44:51 2023 +0000

     cpufreq: amd-pstate-ut: Remove module parameter access

     [ Upstream commit 8d6e5e8268e89979d86501dbb8385ce2e6154de1 ]

     In amd-pstate-ut, shared memory-based systems call
     get_shared_mem() as part of amd_pstate_ut_check_enabled()
     function. This function was written when CONFIG_X86_AMD_PSTATE
     was tristate config and amd_pstate can be built as a module.

     Currently CONFIG_X86_AMD_PSTATE is a boolean config and module
     parameter shared_mem is removed. But amd-pstate-ut code still
     accesses this module parameter. Remove those accesses.

     Fixes: 456ca88d8a52 ("cpufreq: amd-pstate: change amd-pstate driver to be built-in
type")
     Reviewed-by: Mario Limonciello <mario.limonciello@amd.com>
     Reviewed-by: Meng Li <li.meng@amd.com>
     Reviewed-by: Wyes Karny <wyes.karny@amd.com>
     Suggested-by: Wyes Karny <wyes.karny@amd.com>
     Signed-off-by: Swapnil Sapkal <swapnil.sapkal@amd.com>
     [ rjw: Subject edits ]
     Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c9091e03103645a4509d0b4597b3a76dcedec631
Author: Peng Fan <peng.fan@nxp.com>
Date:   Wed Jul 19 09:16:36 2023 +0800

     thermal/of: Fix potential uninitialized value access

     [ Upstream commit f96801f0cfcefc0a16b146596577c53c75ee9773 ]
```

```
    If of_parse_phandle_with_args() called from __thermal_of_bind() or
    __thermal_of_unbind() fails, cooling_spec.np will not be initialized,
    so move the of_node_put() calls below the respective return value checks
    to avoid dereferencing an uninitialized pointer.

    Fixes: 3fd6d6e2b4e8 ("thermal/of: Rework the thermal device tree initialization")
    Signed-off-by: Peng Fan <peng.fan@nxp.com>
    [ rjw: Subject and changelog edits ]
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 61e61e0e0a6ec5d9e996cb98b3d489b5ff7a9815
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Fri Aug 18 14:40:04 2023 -0500

    ACPI: x86: s2idle: Fix a logic error parsing AMD constraints table

    [ Upstream commit 9cc8cd086f05d9a01026c65c98da88561e9c619e ]

    The constraints table should be resetting the `list` object
    after running through all of `info_obj` iterations.

    This adjusts whitespace as well as less code will now be included
    with each loop. This fixes a functional problem is fixed where a
    badly formed package in the inner loop may have incorrect data.

    Fixes: 146f1ed852a8 ("ACPI: PM: s2idle: Add AMD support to handle _DSM")
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 94c476a60af5579e75d1bbca97d8e6cb01eae5a5
Author: Mario Limonciello <mario.limonciello@amd.com>
Date:   Fri Aug 18 14:40:02 2023 -0500

    ACPI: x86: s2idle: Post-increment variables when getting constraints

    [ Upstream commit 3c6b1212d20bbbffcad5709ab0f2d5ed9b5859a8 ]

    When code uses a pre-increment it makes the reader question "why".
    In the constraint fetching code there is no reason for the variables
    to be pre-incremented so adjust to post-increment.
    No intended functional changes.

    Reviewed-by: Kuppuswamy Sathyanarayanan <sathyanarayanan.kuppuswamy@linux.intel.com>
    Suggested-by: Bjorn Helgaas <helgaas@kernel.org>
    Signed-off-by: Mario Limonciello <mario.limonciello@amd.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Stable-dep-of: 9cc8cd086f05 ("ACPI: x86: s2idle: Fix a logic error parsing AMD
constraints table")
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a3baaf5ec281a155b866c985656df1aa6e0dea5f
Author: Bibo Mao <maobibo@loongson.cn>
Date:   Fri Aug 11 17:58:04 2023 +0800

    irqchip/loongson-eiointc: Fix return value checking of eiointc_index

    [ Upstream commit 2e99b73afde18853754c5fae8e8d1a66fe5e3f64 ]

    Return value of function eiointc_index is int, however it is converted
    into uint32_t and then compared smaller than zero, this will cause logic
    problem.
```

```
    Fixes: dd281e1a1a93 ("irqchip: Add Loongson Extended I/O interrupt controller support")
    Signed-off-by: Bibo Mao <maobibo@loongson.cn>
    Reviewed-by: Philippe Mathieu-Daudé <philmd@linaro.org>
    Signed-off-by: Marc Zyngier <maz@kernel.org>
    Link: https://lore.kernel.org/r/20230811095805.2974722-2-maobibo@loongson.cn
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit aadea03a74ab6b5ce9623fa542b23ce0c8995f80
Author: Holger Dengler <dengler@linux.ibm.com>
Date:   Wed Aug 9 14:23:45 2023 +0200

    s390/paes: fix PKEY_TYPE_EP11_AES handling for secure keyblobs

    [ Upstream commit cba33db3fc4dbf2e54294b0e499d2335a3a00d78 ]

    Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
    private keys")' introduced PKEY_TYPE_EP11_AES securekey blobs as a
    supplement to the PKEY_TYPE_EP11 (which won't work in environments
    with session-bound keys). This new keyblobs has a different maximum
    size, so fix paes crypto module to accept also these larger keyblobs.

    Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
    Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
    Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 979afaffcc4739130df4278e55501dd6a629f8d2
Author: Holger Dengler <dengler@linux.ibm.com>
Date:   Fri Aug 4 16:02:58 2023 +0200

    s390/pkey: fix PKEY_TYPE_EP11_AES handling for sysfs attributes

    [ Upstream commit b9352e4b9b9eff949bcc6907b8569b3a1d992f1e ]

    Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
    private keys")' introduced a new PKEY_TYPE_EP11_AES securekey type as
    a supplement to the existing PKEY_TYPE_EP11 (which won't work in
    environments with session-bound keys). The pkey EP11 securekey
    attributes use PKEY_TYPE_EP11_AES (instead of PKEY_TYPE_EP11)
    keyblobs, to make the generated keyblobs usable also in environments,
    where session-bound keys are required.

    There should be no negative impacts to userspace because the internal
    structure of the keyblobs is opaque. The increased size of the
    generated keyblobs is reflected by the changed size of the attributes.

    Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
    Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
    Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5f3501c3145b2a306b91ad226590eedd7c101923
Author: Holger Dengler <dengler@linux.ibm.com>
Date:   Tue Jul 25 13:05:36 2023 +0200

    s390/pkey: fix PKEY_TYPE_EP11_AES handling in PKEY_VERIFYKEY2 IOCTL

    [ Upstream commit 745742dbca11a1b63684ec7032a81aaedcf51fb0 ]

    Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
    private keys")' introduced a new PKEY_TYPE_EP11_AES type for the
    PKEY_VERIFYKEY2 IOCTL to verify keyblobs of this type. Unfortunately,
    all PKEY_VERIFYKEY2 IOCTL requests with keyblobs of this type return
```

```
    with an error (-EINVAL). Fix PKEY_TYPE_EP11_AES handling in
    PKEY_VERIFYKEY2 IOCTL, so that userspace can verify keyblobs of this
    type.

    Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
    Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
    Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 1c7ceb26d51dc3db797f59a17b5d02f5cc0ba0e6
Author: Holger Dengler <dengler@linux.ibm.com>
Date:   Wed Jul 26 16:22:19 2023 +0200

    s390/pkey: fix PKEY_TYPE_EP11_AES handling in PKEY_KBLOB2PROTK[23]

    [ Upstream commit d1fdfb0b2f339cf882c0b5431084a1950b8b73b9 ]

    Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
    private keys")' introduced a new PKEY_TYPE_EP11_AES type for the
    PKEY_KBLOB2PROTK2 and a new IOCTL, PKEY_KBLOB2PROTK3, which both
    allows userspace to convert opaque securekey blobs of this type into
    protectedkey blobs. Unfortunately, all PKEY_KBLOB2PROTK2 and
    PKEY_KBLOB2PROTK3 IOCTL requests with this keyblobs of this type
    return with an error (-EINVAL). Fix PKEY_TYPE_EP11_AES handling in
    PKEY_KBLOB2PROTK2 and PKEY_KBLOB2PROTK3 IOCTLs, so that userspace can
    convert PKEY_TYPE_EP11_AES keyblobs into protectedkey blobs.

    Add a helper function to decode the start and size of the internal
    header as well as start and size of the keyblob payload of an existing
    keyblob. Also validate the length of header and keyblob, as well as
    the keyblob magic.

    Introduce another helper function, which handles a raw key wrapping
    request and do the keyblob decoding in the calling function. Remove
    all other header-related calculations.

    Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
    Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
    Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 542b250d697dd5df3042458bca4e4cc463fec523
Author: Holger Dengler <dengler@linux.ibm.com>
Date:   Tue Jul 25 11:24:47 2023 +0200

    s390/pkey: fix PKEY_TYPE_EP11_AES handling in PKEY_CLR2SECK2 IOCTL

    [ Upstream commit da2863f15945de100b95c72d5656541d30956c5d ]

    Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
    private keys")' introduced PKEY_TYPE_EP11_AES for the PKEY_CLR2SECK2
    IOCTL to convert an AES clearkey into a securekey of this type.
    Unfortunately, all PKEY_CLR2SECK2 IOCTL requests with type
    PKEY_TYPE_EP11_AES return with an error (-EINVAL). Fix the handling
    for PKEY_TYPE_EP11_AES in PKEY_CLR2SECK2 IOCTL, so that userspace can
    convert clearkey blobs into PKEY_TYPE_EP11_AES securekey blobs.

    Cc: stable@vger.kernel.org # v5.10+
    Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
    Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
    Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
    Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
    Stable-dep-of: d1fdfb0b2f33 ("s390/pkey: fix PKEY_TYPE_EP11_AES handling in
```

```
  PKEY_KBLOB2PROTK[23]")
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 5e17fdf876487138fcdfd1b4baacbd92853bafe7
  Author: Holger Dengler <dengler@linux.ibm.com>
  Date:   Tue Jul 25 09:49:55 2023 +0200

      s390/pkey: fix PKEY_TYPE_EP11_AES handling in PKEY_GENSECK2 IOCTL

      [ Upstream commit fb249ce7f7bfd8621a38e4ad401ba74b680786d4 ]

      Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
      private keys")' introduced PKEY_TYPE_EP11_AES for the PKEY_GENSECK2
      IOCTL, to enable userspace to generate securekey blobs of this
      type. Unfortunately, all PKEY_GENSECK2 IOCTL requests for
      PKEY_TYPE_EP11_AES return with an error (-EINVAL). Fix the handling
      for PKEY_TYPE_EP11_AES in PKEY_GENSECK2 IOCTL, so that userspace can
      generate securekey blobs of this type.

      The start of the header and the keyblob, as well as the length need
      special handling, depending on the internal keyversion. Add a helper
      function that splits an uninitialized buffer into start and size of
      the header as well as start and size of the payload, depending on the
      requested keyversion.

      Do the header-related calculations and the raw genkey request handling
      in separate functions. Use the raw genkey request function for
      internal purposes.

      Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
      Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
      Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
      Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 5334225a1a388ab2ce5d6765efcbeaf752f8fac6
  Author: Holger Dengler <dengler@linux.ibm.com>
  Date:   Wed Jul 26 11:33:45 2023 +0200

      s390/pkey: fix/harmonize internal keyblob headers

      [ Upstream commit 37a08f010b7c423b5e4c9ed3b187d21166553007 ]

      Commit 'fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC
      private keys")' introduced PKEY_TYPE_EP11_AES as a supplement to
      PKEY_TYPE_EP11. All pkeys have an internal header/payload structure,
      which is opaque to the userspace. The header structures for
      PKEY_TYPE_EP11 and PKEY_TYPE_EP11_AES are nearly identical and there
      is no reason, why different structures are used. In preparation to fix
      the keyversion handling in the broken PKEY IOCTLs, the same header
      structure is used for PKEY_TYPE_EP11 and PKEY_TYPE_EP11_AES. This
      reduces the number of different code paths and increases the
      readability.

      Fixes: fa6999e326fe ("s390/pkey: support CCA and EP11 secure ECC private keys")
      Signed-off-by: Holger Dengler <dengler@linux.ibm.com>
      Reviewed-by: Ingo Franzki <ifranzki@linux.ibm.com>
      Signed-off-by: Heiko Carstens <hca@linux.ibm.com>
      Signed-off-by: Sasha Levin <sashal@kernel.org>

  commit 22e3f7a9221e77f9ad796544714482f3f6b9c300
  Author: Nysal Jan K.A <nysal@linux.ibm.com>
  Date:   Mon Aug 14 13:39:27 2023 +0530

      selftests/futex: Order calls to futex_lock_pi
```

```
    [ Upstream commit fbf4dec702774286db409815ffb077711a96b824 ]

    Observed occassional failures in the futex_wait_timeout test:

    ok 1 futex_wait relative succeeds
    ok 2 futex_wait_bitset realtime succeeds
    ok 3 futex_wait_bitset monotonic succeeds
    ok 4 futex_wait_requeue_pi realtime succeeds
    ok 5 futex_wait_requeue_pi monotonic succeeds
    not ok 6 futex_lock_pi realtime returned 0
    ......

    The test expects the child thread to complete some steps before
    the parent thread gets to run. There is an implicit expectation
    of the order of invocation of futex_lock_pi between the child thread
    and the parent thread. Make this order explicit. If the order is
    not met, the futex_lock_pi call in the parent thread succeeds and
    will not timeout.

    Fixes: f4addd54b161 ("selftests: futex: Expand timeout test")
    Signed-off-by: Nysal Jan K.A <nysal@linux.ibm.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 5b9836417c99081de40bf76c23cd229b7a001527
Author: Xu Yang <xu.yang_2@nxp.com>
Date:   Fri Aug 11 09:54:38 2023 +0800

    perf/imx_ddr: don't enable counter0 if none of 4 counters are used

    [ Upstream commit f4e2bd91ddf5e8543cbe7ad80b3fba3d2dc63fa3 ]

    In current driver, counter0 will be enabled after ddr_perf_pmu_enable()
    is called even though none of the 4 counters are used. This will cause
    counter0 continue to count until ddr_perf_pmu_disabled() is called. If
    pmu is not disabled all the time, the pmu interrupt will be asserted
    from time to time due to counter0 will overflow and irq handler will
    clear it. It's not an expected behavior. This patch will not enable
    counter0 if none of 4 counters are used.

    Fixes: 9a66d36cc7ac ("drivers/perf: imx_ddr: Add DDR performance counter support to
perf")
    Signed-off-by: Xu Yang <xu.yang_2@nxp.com>
    Reviewed-by: Frank Li <Frank.Li@nxp.com>
    Link: https://lore.kernel.org/r/20230811015438.1999307-2-xu.yang_2@nxp.com
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 843d4fc3a5ac0cab44a8388c3534cdee45e9b518
Author: Cyril Hrubis <chrubis@suse.cz>
Date:   Wed Aug 2 17:19:05 2023 +0200

    sched/rt: Fix sysctl_sched_rr_timeslice intial value

    [ Upstream commit c7fcb99877f9f542c918509b2801065adcaf46fa ]

    There is a 10% rounding error in the intial value of the
    sysctl_sched_rr_timeslice with CONFIG_HZ_300=y.

    This was found with LTP test sched_rr_get_interval01:

    sched_rr_get_interval01.c:57: TPASS: sched_rr_get_interval() passed
    sched_rr_get_interval01.c:64: TPASS: Time quantum 0s 99999990ns
    sched_rr_get_interval01.c:72: TFAIL: /proc/sys/kernel/sched_rr_timeslice_ms != 100 got
```

```
 90
     sched_rr_get_interval01.c:57: TPASS: sched_rr_get_interval() passed
     sched_rr_get_interval01.c:64: TPASS: Time quantum 0s 99999990ns
     sched_rr_get_interval01.c:72: TFAIL: /proc/sys/kernel/sched_rr_timeslice_ms != 100 got
 90

     What this test does is to compare the return value from the
     sched_rr_get_interval() and the sched_rr_timeslice_ms sysctl file and
     fails if they do not match.

     The problem it found is the intial sysctl file value which was computed as:

     static int sysctl_sched_rr_timeslice = (MSEC_PER_SEC / HZ) * RR_TIMESLICE;

     which works fine as long as MSEC_PER_SEC is multiple of HZ, however it
     introduces 10% rounding error for CONFIG_HZ_300:

     (MSEC_PER_SEC / HZ) * (100 * HZ / 1000)

     (1000 / 300) * (100 * 300 / 1000)

     3 * 30 = 90

     This can be easily fixed by reversing the order of the multiplication
     and division. After this fix we get:

     (MSEC_PER_SEC * (100 * HZ / 1000)) / HZ

     (1000 * (100 * 300 / 1000)) / 300

     (1000 * 30) / 300 = 100

     Fixes: 975e155ed873 ("sched/rt: Show the 'sched_rr_timeslice' SCHED_RR timeslice tuning
 knob in milliseconds")
     Signed-off-by: Cyril Hrubis <chrubis@suse.cz>
     Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
     Reviewed-by: Petr Vorel <pvorel@suse.cz>
     Acked-by: Mel Gorman <mgorman@suse.de>
     Tested-by: Petr Vorel <pvorel@suse.cz>
     Link: https://lore.kernel.org/r/20230802151906.25258-2-chrubis@suse.cz
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 13bbf4bf1d3337fe717905e29a525b4f9c513fae
 Author: Mark Brown <broonie@kernel.org>
 Date:   Mon Jul 31 14:58:48 2023 +0100

     arm64/fpsimd: Only provide the length to cpufeature for xCR registers

     [ Upstream commit 01948b09edc3fecf8486c57c2d2fb8b80886f3d0 ]

     For both SVE and SME we abuse the generic register field comparison
     support in the cpufeature code as part of our detection of unsupported
     variations in the vector lengths available to PEs, reporting the maximum
     vector lengths via ZCR_EL1.LEN and SMCR_EL1.LEN.  Since these are
     configuration registers rather than identification registers the
     assumptions the cpufeature code makes about how unknown bitfields behave
     are invalid, leading to warnings when SME features like FA64 are enabled
     and we hotplug a CPU:

       CPU features: SANITY CHECK: Unexpected variation in SYS_SMCR_EL1. Boot CPU:
 0x0000000000000f, CPU3: 0x0000008000000f
       CPU features: Unsupported CPU feature variation detected.

     SVE has no controls other than the vector length so is not yet impacted
     but the same issue will apply there if any are defined.
```

```
    Since the only field we are interested in having the cpufeature code
    handle is the length field and we use a custom read function to obtain
    the value we can avoid these warnings by filtering out all other bits
    when we return the register value, if we're doing that we don't need to
    bother reading the register at all and can simply use the RDVL/RDSVL
    value we were filling in instead.

    Fixes: 2e0f2478ea37 ("arm64/sve: Probe SVE capabilities and usable vector lengths")
    FixeS: b42990d3bf77 ("arm64/sme: Identify supported SME vector lengths at boot")
    Signed-off-by: Mark Brown <broonie@kernel.org>
    Reviewed-by: Catalin Marinas <catalin.marinas@arm.com>
    Link: https://lore.kernel.org/r/20230731-arm64-sme-fa64-hotplug-v2-1-
7714c00dd902@kernel.org
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a355b274052c3b53ff29a18eead103b8911f565e
Author: Ard Biesheuvel <ardb@kernel.org>
Date:   Mon Aug 7 18:26:58 2023 +0200

    x86/decompressor: Don't rely on upper 32 bits of GPRs being preserved

    [ Upstream commit 264b82fdb4989cf6a44a2bcd0c6ea05e8026b2ac ]

    The 4-to-5 level mode switch trampoline disables long mode and paging in
    order to be able to flick the LA57 bit. According to section 3.4.1.1 of
    the x86 architecture manual [0], 64-bit GPRs might not retain the upper
    32 bits of their contents across such a mode switch.

    Given that RBP, RBX and RSI are live at this point, preserve them on the
    stack, along with the return address that might be above 4G as well.

    [0] Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1: Basic
Architecture

      "Because the upper 32 bits of 64-bit general-purpose registers are
       undefined in 32-bit modes, the upper 32 bits of any general-purpose
       register are not preserved when switching from 64-bit mode to a 32-bit
       mode (to protected mode or compatibility mode). Software must not
       depend on these bits to maintain a value after a 64-bit to 32-bit
       mode switch."

    Fixes: 194a9749c73d650c ("x86/boot/compressed/64: Handle 5-level paging boot if kernel
is above 4G")
    Signed-off-by: Ard Biesheuvel <ardb@kernel.org>
    Signed-off-by: Borislav Petkov (AMD) <bp@alien8.de>
    Link: https://lore.kernel.org/r/20230807162720.545787-2-ardb@kernel.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 25199be8cadc1992e8aca0b5892dbecc69807cca
Author: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
Date:   Mon Jul 31 20:56:35 2023 +0200

    cpuidle: teo: Update idle duration estimate when choosing shallower state

    [ Upstream commit 3f0b0966b30982e843950b170b7a9ddfd8094428 ]

    The TEO governor takes CPU utilization into account by refining idle state
    selection when the utilization is above a certain threshold.  This is done by
    choosing an idle state shallower than the previously selected one.

    However, when doing this, the idle duration estimate needs to be
    adjusted so as to prevent the scheduler tick from being stopped when the
    candidate idle state is shallow, which may lead to excessive energy
```

```
    usage if the CPU is not woken up quickly enough going forward.
    Moreover, if the scheduler tick has been stopped already and the new
    idle duration estimate is too small, the replacement candidate state
    cannot be used.

    Modify the relevant code to take the above observations into account.

    Fixes: 9ce0f7c4bc64 ("cpuidle: teo: Introduce util-awareness")
    Link: https://lore.kernel.org/linux-
pm/CAJZ5v0jJxHj65r2HXBTd3wfbZtsg=_StzwO1kA5STDnaPe_dWA@mail.gmail.com
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Reviewed-and-tested-by: Kajetan Puchalski <kajetan.puchalski@arm.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit cd62464707e11919a0a1017119d1cfb3e4502f53
Author: Randy Dunlap <rdunlap@infradead.org>
Date:   Sun Jul 30 20:07:40 2023 -0700

    sched/psi: Select KERNFS as needed

    [ Upstream commit 98dfdd9ee93995a408192dbbf3dd219ba23e3738 ]

    Users of KERNFS should select it to enforce its being built, so
    do this to prevent a build error.

    In file included from ../kernel/sched/build_utility.c:97:
    ../kernel/sched/psi.c: In function 'psi_trigger_poll':
    ../kernel/sched/psi.c:1479:17: error: implicit declaration of function
'kernfs_generic_poll' [-Werror=implicit-function-declaration]
     1479 |                 kernfs_generic_poll(t->of, wait);

    Fixes: aff037078eca ("sched/psi: use kernfs polling functions for PSI trigger polling")
    Reported-by: kernel test robot <lkp@intel.com>
    Signed-off-by: Randy Dunlap <rdunlap@infradead.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Acked-by: Suren Baghdasaryan <surenb@google.com>
    Link: lore.kernel.org/r/202307310732.r65EQFY0-lkp@intel.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 88a02078f68ae6a11ec43705fb2af259e9644071
Author: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
Date:   Mon Jul 17 19:55:05 2023 +0200

    arm64/ptrace: Clean up error handling path in sve_set_common()

    [ Upstream commit 5f69ca4229c7d8e23f238174827ee7aa49b0bcb2 ]

    All error handling paths go to 'out', except this one. Be consistent and
    also branch to 'out' here.

    Fixes: e12310a0d30f ("arm64/sme: Implement ptrace support for streaming mode SVE
registers")
    Signed-off-by: Christophe JAILLET <christophe.jaillet@wanadoo.fr>
    Reviewed-by: Mark Brown <broonie@kernel.org>
    Reviewed-by: Anshuman Khandual <anshuman.khandual@arm.com>
    Link:
https://lore.kernel.org/r/aa61301ed2dfd079b74b37f7fede5f179ac3087a.1689616473.git.christophe
.jaillet@wanadoo.fr
    Signed-off-by: Will Deacon <will@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 612a064d80d24bafe739f5e0a654b657663dd1d6
Author: Vincent Guittot <vincent.guittot@linaro.org>
Date:   Thu Jul 6 15:51:44 2023 +0200
```

```
    sched/fair: remove util_est boosting

    [ Upstream commit c2e164ac33f75e0acb93004960c73bd9166d3d35 ]

    There is no need to use runnable_avg when estimating util_est and that
    even generates wrong behavior because one includes blocked tasks whereas
    the other one doesn't. This can lead to accounting twice the waking task p,
    once with the blocked runnable_avg and another one when adding its
    util_est.

    cpu's runnable_avg is already used when computing util_avg which is then
    compared with util_est.

    In some situation, feec will not select prev_cpu but another one on the
    same performance domain because of higher max_util

    Fixes: 7d0583cf9ec7 ("sched/fair, cpufreq: Introduce 'runnable boosting'")
    Signed-off-by: Vincent Guittot <vincent.guittot@linaro.org>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Reviewed-by: Dietmar Eggemann <dietmar.eggemann@arm.com>
    Tested-by: Dietmar Eggemann <dietmar.eggemann@arm.com>
    Link: https://lore.kernel.org/r/20230706135144.324311-1-vincent.guittot@linaro.org
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8eb28aebd7e73a517e0e57100bdfd17f2e5e3834
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 16:14:52 2023 +0300

    selftests/resctrl: Close perf value read fd on errors

    [ Upstream commit 51a0c3b7f028169e40db930575dd01fe81c3e765 ]

    Perf event fd (fd_lm) is not closed when run_fill_buf() returns error.

    Close fd_lm only in cat_val() to make it easier to track it is always
    closed.

    Fixes: 790bf585b0ee ("selftests/resctrl: Add Cache Allocation Technology (CAT)
selftest")
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Tested-by: Babu Moger <babu.moger@amd.com>
    Tested-by: Shaopeng Tan (Fujitsu) <tan.shaopeng@fujitsu.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit db8e2eaaeba8d52dc0f4819796feeafd18c7a952
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 16:14:51 2023 +0300

    selftests/resctrl: Unmount resctrl FS if child fails to run benchmark

    [ Upstream commit f99e413eb54652e2436cc56d081176bc9a34cd8d ]

    A child calls PARENT_EXIT() when it fails to run a benchmark to kill
    the parent process. PARENT_EXIT() lacks unmount for the resctrl FS and
    the parent won't be there to unmount it either after it gets killed.

    Add the resctrl FS unmount also to PARENT_EXIT().

    Fixes: 591a6e8588fc ("selftests/resctrl: Add basic resctrl file system operations and
data")
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Reviewed-by: Reinette Chatre <reinette.chatre@intel.com>
    Tested-by: Babu Moger <babu.moger@amd.com>
    Tested-by: Shaopeng Tan (Fujitsu) <tan.shaopeng@fujitsu.com>
```

```
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b14094cd7d1feaee65b1d959aea13a6c53ffd03a
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 16:14:50 2023 +0300

    selftests/resctrl: Don't leak buffer in fill_cache()

    [ Upstream commit 2d320b1029ee7329ee0638181be967789775b962 ]

    The error path in fill_cache() does return before the allocated buffer
    is freed leaking the buffer.

    The leak was introduced when fill_cache_read() started to return errors
    in commit c7b607fa9325 ("selftests/resctrl: Fix null pointer
    dereference on open failed"), before that both fill functions always
    returned 0.

    Move free() earlier to prevent the mem leak.

    Fixes: c7b607fa9325 ("selftests/resctrl: Fix null pointer dereference on open failed")
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Reviewed-by: Reinette Chatre <reinette.chatre@intel.com>
    Tested-by: Babu Moger <babu.moger@amd.com>
    Tested-by: Shaopeng Tan (Fujitsu) <tan.shaopeng@fujitsu.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 0a6055306ec6a65b07d47e66d4257190d9656bab
Author: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
Date:   Mon Jul 17 16:14:49 2023 +0300

    selftests/resctrl: Add resctrl.h into build deps

    [ Upstream commit 8e289f4542890168705219e54f0231dccfabddbe ]

    Makefile only lists *.c as build dependencies for the resctrl_tests
    executable which excludes resctrl.h.

    Add *.h to wildcard() to include resctrl.h.

    Fixes: 591a6e8588fc ("selftests/resctrl: Add basic resctrl file system operations and
data")
    Signed-off-by: Ilpo Järvinen <ilpo.jarvinen@linux.intel.com>
    Reviewed-by: Reinette Chatre <reinette.chatre@intel.com>
    Tested-by: Babu Moger <babu.moger@amd.com>
    Tested-by: Shaopeng Tan (Fujitsu) <tan.shaopeng@fujitsu.com>
    Signed-off-by: Shuah Khan <skhan@linuxfoundation.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 2147ee4774c47a14a54c6e31207a79c4bd17c28a
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Fri Jul 21 18:16:34 2023 +0530

    OPP: Fix passing 0 to PTR_ERR in _opp_attach_genpd()

    [ Upstream commit d920920f85a82c1c806a4143871a0e8f534732f2 ]

    If dev_pm_domain_attach_by_name() returns NULL, then 0 will be passed to
    PTR_ERR() as reported by the smatch warning below:

    drivers/opp/core.c:2456 _opp_attach_genpd() warn: passing zero to 'PTR_ERR'

    Fix it by checking for the non-NULL virt_dev pointer before passing it to
```

```
    PTR_ERR. Otherwise return -ENODEV.

    Fixes: 4ea9496cbc95 ("opp: Fix error check in dev_pm_opp_attach_genpd()")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 25130b27e0352acb83e91c467853eb9afad3b644
Author: Manivannan Sadhasivam <mani@kernel.org>
Date:   Fri Jul 21 18:16:33 2023 +0530

    OPP: Fix potential null ptr dereference in dev_pm_opp_get_required_pstate()

    [ Upstream commit 7ddd8deb1c3c0363a7e14fafb5df26e2089a69a5 ]

    "opp" pointer is dereferenced before the IS_ERR_OR_NULL() check. Fix it by
    removing the dereference to cache opp_table and dereference it directly
    where opp_table is used.

    This fixes the following smatch warning:

    drivers/opp/core.c:232 dev_pm_opp_get_required_pstate() warn: variable
    dereferenced before IS_ERR check 'opp' (see line 230)

    Fixes: 84cb7ff35fcf ("OPP: pstate is only valid for genpd OPP tables")
    Signed-off-by: Manivannan Sadhasivam <manivannan.sadhasivam@linaro.org>
    Signed-off-by: Viresh Kumar <viresh.kumar@linaro.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 8022b64fb7daa6135d9f7b0e2f7b5b8e9e5179c9
Author: Cristian Marussi <cristian.marussi@arm.com>
Date:   Tue Jul 18 11:17:26 2023 +0100

    powercap: arm_scmi: Remove recursion while parsing zones

    [ Upstream commit 3e767d6850f867cc33ac16ca097350a1d2417982 ]

    Powercap zones can be defined as arranged in a hierarchy of trees and when
    registering a zone with powercap_register_zone(), the kernel powercap
    subsystem expects this to happen starting from the root zones down to the
    leaves; on the other side, de-registration by powercap_deregister_zone()
    must begin from the leaf zones.

    Available SCMI powercap zones are retrieved dynamically from the platform
    at probe time and, while any defined hierarchy between the zones is
    described properly in the zones descriptor, the platform returns the
    availables zones with no particular well-defined order: as a consequence,
    the trees possibly composing the hierarchy of zones have to be somehow
    walked properly to register the retrieved zones from the root.

    Currently the ARM SCMI Powercap driver walks the zones using a recursive
    algorithm; this approach, even though correct and tested can lead to kernel
    stack overflow when processing a returned hierarchy of zones composed by
    particularly high trees.

    Avoid possible kernel stack overflow by substituting the recursive approach
    with an iterative one supported by a dynamically allocated stack-like data
    structure.

    Fixes: b55eef5226b7 ("powercap: arm_scmi: Add SCMI Powercap based driver")
    Signed-off-by: Cristian Marussi <cristian.marussi@arm.com>
    Acked-by: Sudeep Holla <sudeep.holla@arm.com>
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit 6e779ff223a8bbdcee4ec5d1e51765ebaa323ac9
Author: Paul E. McKenney <paulmck@kernel.org>
Date:   Wed Jun 7 11:59:49 2023 -0700

    clocksource: Handle negative skews in "skew is too large" messages

    [ Upstream commit e40806e9bcf8aaa86dbf0d484e7cf3cfa09cb86c ]

    The nanosecond-to-millisecond skew computation uses unsigned arithmetic,
    which produces user-unfriendly large positive numbers for negative skews.
    Therefore, use signed arithmetic for this computation in order to preserve
    the negativity.

    Reported-by: Chris Bainbridge <chris.bainbridge@gmail.com>
    Reported-by: Feng Tang <feng.tang@intel.com>
    Fixes: dd029269947a ("clocksource: Improve "skew is too large" messages")
    Reviewed-by: Feng Tang <feng.tang@intel.com>
    Tested-by: Chris Bainbridge <chris.bainbridge@gmail.com>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 70a2856fd1d0a040c876ba9e3f89b949ae92e4dd
Author: Waiman Long <longman@redhat.com>
Date:   Fri Jul 7 13:53:55 2023 -0400

    refscale: Fix uninitalized use of wait_queue_head_t

    [ Upstream commit f5063e8948dad7f31adb007284a5d5038ae31bb8 ]

    Running the refscale test occasionally crashes the kernel with the
    following error:

    [ 8569.952896] BUG: unable to handle page fault for address: fffffffffffffe8
    [ 8569.952900] #PF: supervisor read access in kernel mode
    [ 8569.952902] #PF: error_code(0x0000) - not-present page
    [ 8569.952904] PGD c4b048067 P4D c4b049067 PUD c4b04b067 PMD 0
    [ 8569.952910] Oops: 0000 [#1] PREEMPT_RT SMP NOPTI
    [ 8569.952916] Hardware name: Dell Inc. PowerEdge R750/0WMWCR, BIOS 1.2.4 05/28/2021
    [ 8569.952917] RIP: 0010:prepare_to_wait_event+0x101/0x190
      :
    [ 8569.952940] Call Trace:
    [ 8569.952941]  <TASK>
    [ 8569.952944]  ref_scale_reader+0x380/0x4a0 [refscale]
    [ 8569.952959]  kthread+0x10e/0x130
    [ 8569.952966]  ret_from_fork+0x1f/0x30
    [ 8569.952973]  </TASK>

    The likely cause is that init_waitqueue_head() is called after the call to
    the torture_create_kthread() function that creates the ref_scale_reader
    kthread.  Although this init_waitqueue_head() call will very likely
    complete before this kthread is created and starts running, it is
    possible that the calling kthread will be delayed between the calls to
    torture_create_kthread() and init_waitqueue_head().  In this case, the
    new kthread will use the waitqueue head before it is properly initialized,
    which is not good for the kernel's health and well-being.

    The above crash happened here:

            static inline void __add_wait_queue(...)
            {
                    :
                    if (!(wq->flags & WQ_FLAG_PRIORITY)) <=== Crash here

    The offset of flags from list_head entry in wait_queue_entry is
    -0x18. If reader_tasks[i].wq.head.next is NULL as allocated reader_task
```

```
    structure is zero initialized, the instruction will try to access address
    0xfffffffffffffe8, which is exactly the fault address listed above.

    This commit therefore invokes init_waitqueue_head() before creating
    the kthread.

    Fixes: 653ed64b01dc ("refperf: Add a test to measure performance of read-side
synchronization")
    Signed-off-by: Waiman Long <longman@redhat.com>
    Reviewed-by: Qiuxu Zhuo <qiuxu.zhuo@intel.com>
    Reviewed-by: Davidlohr Bueso <dave@stgolabs.net>
    Acked-by: Joel Fernandes (Google) <joel@joelfernandes.org>
    Signed-off-by: Paul E. McKenney <paulmck@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit c217d8f2e1c3ebe6ddefa1b226e33e6b54aa288d
Author: Uros Bizjak <ubizjak@gmail.com>
Date:   Sat Jul 8 11:00:36 2023 +0200

    locking/arch: Avoid variable shadowing in local_try_cmpxchg()

    [ Upstream commit d6b45484c130f4095313ae3edeb4aae662c12fb1 ]

    Several architectures define arch_try_local_cmpxchg macro using
    internal temporary variables named ___old, __old or _old. Remove
    temporary varible in local_try_cmpxchg to avoid variable shadowing.

    No functional change intended.

    Fixes: d994f2c8e241 ("locking/arch: Wire up local_try_cmpxchg()")
    Closes: https://lore.kernel.org/lkml/CAFGhKbyxtuk=LoW-E3yLXgcmR93m+Dfo5-
u9oQA_YC5Fcy_t9g@mail.gmail.com/
    Reported-by: Charlemagne Lasse <charlemagnelasse@gmail.com>
    Signed-off-by: Uros Bizjak <ubizjak@gmail.com>
    Signed-off-by: Peter Zijlstra (Intel) <peterz@infradead.org>
    Link: https://lkml.kernel.org/r/20230708090048.63046-1-ubizjak@gmail.com
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 9f9e7fea0ff2ad237ceda88597d2a5936d67c4f0
Author: Jingbo Xu <jefflexu@linux.alibaba.com>
Date:   Tue Aug 22 19:05:30 2023 +0800

    erofs: release ztailpacking pclusters properly

    [ Upstream commit 91b1ad0815fbb1095c8b9e8a2bf4201186afe304 ]

    Currently ztailpacking pclusters are chained with FOLLOWED_NOINPLACE and
    not recorded into the managed_pslots XArray.

    After commit 7674a42f35ea ("erofs: use struct lockref to replace
    handcrafted approach"), ztailpacking pclusters won't be freed with
    erofs_workgroup_put() anymore, which will cause the following issue:

    BUG erofs_pcluster-1 (Tainted: G           OE     ): Objects remaining in
erofs_pcluster-1 on __kmem_cache_shutdown()

    Use z_erofs_free_pcluster() directly to free ztailpacking pclusters.

    Fixes: 7674a42f35ea ("erofs: use struct lockref to replace handcrafted approach")
    Signed-off-by: Jingbo Xu <jefflexu@linux.alibaba.com>
    Reviewed-by: Gao Xiang <hsiangkao@linux.alibaba.com>
    Reviewed-by: Chao Yu <chao@kernel.org>
    Link: https://lore.kernel.org/r/20230822110530.96831-1-jefflexu@linux.alibaba.com
    Signed-off-by: Gao Xiang <hsiangkao@linux.alibaba.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>
```

```
commit d3cfa44164688a076e8b476cafb5df87d07cfa63
Author: Naohiro Aota <naohiro.aota@wdc.com>
Date:   Sat Aug 19 01:26:07 2023 +0900

    btrfs: zoned: skip splitting and logical rewriting on pre-alloc write

    [ Upstream commit c02d35d89b317994bd713ba82e160c5e7f22d9c8 ]

    When doing a relocation, there is a chance that at the time of
    btrfs_reloc_clone_csums(), there is no checksum for the corresponding
    region.

    In this case, btrfs_finish_ordered_zoned()'s sum points to an invalid item
    and so ordered_extent's logical is set to some invalid value. Then,
    btrfs_lookup_block_group() in btrfs_zone_finish_endio() failed to find a
    block group and will hit an assert or a null pointer dereference as
    following.

    This can be reprodcued by running btrfs/028 several times (e.g, 4 to 16
    times) with a null_blk setup. The device's zone size and capacity is set to
    32 MB and the storage size is set to 5 GB on my setup.

        KASAN: null-ptr-deref in range [0x0000000000000088-0x000000000000008f]
        CPU: 6 PID: 3105720 Comm: kworker/u16:13 Tainted: G        W          6.5.0-rc6-kts+
#1
        Hardware name: Supermicro Super Server/X10SRL-F, BIOS 2.0 12/17/2015
        Workqueue: btrfs-endio-write btrfs_work_helper [btrfs]
        RIP: 0010:btrfs_zone_finish_endio.part.0+0x34/0x160 [btrfs]
        Code: 41 54 49 89 fc 55 48 89 f5 53 e8 57 7d fc ff 48 8d b8 88 00 00 00 48 89 c3 48
b8 00 00 00 00 00
        > 3c 02 00 0f 85 02 01 00 00 f6 83 88 00 00 00 01 0f 84 a8 00 00
        RSP: 0018:ffff88833cf87b08 EFLAGS: 00010206
        RAX: dfffc0000000000 RBX: 0000000000000000 RCX: 0000000000000000
        RDX: 0000000000000011 RSI: 0000000000000004 RDI: 0000000000000088
        RBP: 0000000000000002 R08: 0000000000000001 R09: ffffed102877b827
        R10: ffff888143bdc13b R11: ffff888125b1cbc0 R12: ffff888143bdc000
        R13: 0000000000007000 R14: ffff888125b1cba8 R15: 0000000000000000
        FS:  0000000000000000(0000) GS:ffff88881e500000(0000) knlGS:0000000000000000
        CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
        CR2: 00007f3ed85223d5 CR3: 00000001519b4005 CR4: 00000000001706e0
        Call Trace:
         <TASK>
         ? die_addr+0x3c/0xa0
         ? exc_general_protection+0x148/0x220
         ? asm_exc_general_protection+0x22/0x30
         ? btrfs_zone_finish_endio.part.0+0x34/0x160 [btrfs]
         ? btrfs_zone_finish_endio.part.0+0x19/0x160 [btrfs]
         btrfs_finish_one_ordered+0x7b8/0x1de0 [btrfs]
         ? rcu_is_watching+0x11/0xb0
         ? lock_release+0x47a/0x620
         ? btrfs_finish_ordered_zoned+0x59b/0x800 [btrfs]
         ? __pfx_btrfs_finish_one_ordered+0x10/0x10 [btrfs]
         ? btrfs_finish_ordered_zoned+0x358/0x800 [btrfs]
         ? __smp_call_single_queue+0x124/0x350
         ? rcu_is_watching+0x11/0xb0
         btrfs_work_helper+0x19f/0xc60 [btrfs]
         ? __pfx_try_to_wake_up+0x10/0x10
         ? _raw_spin_unlock_irq+0x24/0x50
         ? rcu_is_watching+0x11/0xb0
         process_one_work+0x8c1/0x1430
         ? __pfx_lock_acquire+0x10/0x10
         ? __pfx_process_one_work+0x10/0x10
         ? __pfx_do_raw_spin_lock+0x10/0x10
         ? _raw_spin_lock_irq+0x52/0x60
```

```
          worker_thread+0x100/0x12c0
          ? __kthread_parkme+0xc1/0x1f0
          ? __pfx_worker_thread+0x10/0x10
          kthread+0x2ea/0x3c0
          ? __pfx_kthread+0x10/0x10
          ret_from_fork+0x30/0x70
          ? __pfx_kthread+0x10/0x10
          ret_from_fork_asm+0x1b/0x30
          </TASK>
```

    On the zoned mode, writing to pre-allocated region means data relocation
    write. Such write always uses WRITE command so there is no need of splitting
    and rewriting logical address. Thus, we can just skip the function for the
    case.

    Fixes: cbfce4c7fbde ("btrfs: optimize the logical to physical mapping for zoned writes")
    Signed-off-by: Naohiro Aota <naohiro.aota@wdc.com>
    Signed-off-by: David Sterba <dsterba@suse.com>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e8e4d22e3ab72464db61dfc4982335825fdf353d
Author: Kees Cook <keescook@chromium.org>
Date: Thu Aug 10 12:54:19 2023 -0700

    ARM: ptrace: Restore syscall skipping for tracers

    [Upstream commit 4697b5848bd933f68ebd04836362c8de0cacaf71]

    Since commit 4e57a4ddf6b0 ("ARM: 9107/1: syscall: always store
    thread_info->abi_syscall"), the seccomp selftests "syscall_errno"
    and "syscall_faked" have been broken. Both seccomp and PTRACE depend
    on using the special value of "-1" for skipping syscalls. This value
    wasn't working because it was getting masked by __NR_SYSCALL_MASK in
    both PTRACE_SET_SYSCALL and get_syscall_nr().

    Explicitly test for -1 in PTRACE_SET_SYSCALL and get_syscall_nr(),
    leaving it exposed when present, allowing tracers to skip syscalls
    again.

    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Arnd Bergmann <arnd@kernel.org>
    Cc: Lecopzer Chen <lecopzer.chen@mediatek.com>
    Cc: Oleg Nesterov <oleg@redhat.com>
    Cc: linux-arm-kernel@lists.infradead.org
    Fixes: 4e57a4ddf6b0 ("ARM: 9107/1: syscall: always store thread_info->abi_syscall")
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20230810195422.2304827-2-keescook@chromium.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit f8e97a03685b4d7256fe333676d81194051fba6a
Author: Kees Cook <keescook@chromium.org>
Date: Thu Aug 10 12:54:18 2023 -0700

    ARM: ptrace: Restore syscall restart tracing

    [Upstream commit cf007647475b5090819c5fe8da771073145c7334]

    Since commit 4e57a4ddf6b0 ("ARM: 9107/1: syscall: always store
    thread_info->abi_syscall"), the seccomp selftests "syscall_restart" has
    been broken. This was caused by the restart syscall not being stored to
    "abi_syscall" during restart setup before branching to the "local_restart"
    label. Tracers would see the wrong syscall, and scno would get overwritten
    while returning from the TIF_WORK path. Add the missing store.

```
    Cc: Russell King <linux@armlinux.org.uk>
    Cc: Arnd Bergmann <arnd@kernel.org>
    Cc: Lecopzer Chen <lecopzer.chen@mediatek.com>
    Cc: Oleg Nesterov <oleg@redhat.com>
    Cc: linux-arm-kernel@lists.infradead.org
    Fixes: 4e57a4ddf6b0 ("ARM: 9107/1: syscall: always store thread_info->abi_syscall")
    Reviewed-by: Arnd Bergmann <arnd@arndb.de>
    Link: https://lore.kernel.org/r/20230810195422.2304827-1-keescook@chromium.org
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit a8d3a6e285f6bc92b37ffd30028869642b1d2682
Author: David Howells <dhowells@redhat.com>
Date:   Tue Aug 8 07:34:20 2023 -0400

    vfs, security: Fix automount superblock LSM init problem, preventing NFS sb sharing

    [ Upstream commit d80a8f1b58c2bc8d7c6bfb65401ea4f7ec8cddc2 ]

    When NFS superblocks are created by automounting, their LSM parameters
    aren't set in the fs_context struct prior to sget_fc() being called,
    leading to failure to match existing superblocks.

    This bug leads to messages like the following appearing in dmesg when
    fscache is enabled:

        NFS: Cache volume key already in use
(nfs,4.2,2,108,106a8c0,1,,,,100000,100000,2ee,3a98,1d4c,3a98,1)

    Fix this by adding a new LSM hook to load fc->security for submount
    creation.

    Signed-off-by: David Howells <dhowells@redhat.com>
    Signed-off-by: Jeff Layton <jlayton@kernel.org>
    Link:
https://lore.kernel.org/r/165962680944.3334508.6610023900349142034.stgit@warthog.procyon.org
.uk/ # v1
    Link:
https://lore.kernel.org/r/165962729225.3357250.14350728846471527137.stgit@warthog.procyon.or
g.uk/ # v2
    Link:
https://lore.kernel.org/r/165970659095.2812394.6868894171102318796.stgit@warthog.procyon.org
.uk/ # v3
    Link:
https://lore.kernel.org/r/166133579016.3678898.6283195019480567275.stgit@warthog.procyon.org
.uk/ # v4
    Link: https://lore.kernel.org/r/217595.1662033775@warthog.procyon.org.uk/ # v5
    Fixes: 9bc61ab18b1d ("vfs: Introduce fs_context, switch vfs_kern_mount() to it.")
    Fixes: 779df6a5480f ("NFS: Ensure security label is set for root inode")
    Tested-by: Jeff Layton <jlayton@kernel.org>
    Acked-by: Casey Schaufler <casey@schaufler-ca.com>
    Acked-by: "Christian Brauner (Microsoft)" <brauner@kernel.org>
    Acked-by: Paul Moore <paul@paul-moore.com>
    Reviewed-by: Jeff Layton <jlayton@kernel.org>
    Message-Id: <20230808-master-v9-1-e0ecde888221@kernel.org>
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit b3547e212afd6a8e62ad01e41e0996dc60f72c26
Author: Kees Cook <keescook@chromium.org>
Date:   Mon Aug 7 10:43:58 2023 -0700

    selftests/harness: Actually report SKIP for signal tests

    [ Upstream commit b3d46e11fec0c5a8972e5061bb1462119ae5736d ]
```

```
    Tests that were expecting a signal were not correctly checking for a
    SKIP condition. Move the check before the signal checking when
    processing test result.

    Cc: Shuah Khan <shuah@kernel.org>
    Cc: Andy Lutomirski <luto@amacapital.net>
    Cc: Will Drewry <wad@chromium.org>
    Cc: linux-kselftest@vger.kernel.org
    Fixes: 9847d24af95c ("selftests/harness: Refactor XFAIL into SKIP")
    Signed-off-by: Kees Cook <keescook@chromium.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit d4ed5bf06257d6dfe3aee538be5855d987d31bdf
Author: Christian Brauner <brauner@kernel.org>
Date:   Tue Aug 1 18:17:04 2023 +0200

    tmpfs: verify {g,u}id mount options correctly

    [ Upstream commit 0200679fc7953177941e41c2a4241d0b6c2c5de8 ]

    A while ago we received the following report:

    "The other outstanding issue I noticed comes from the fact that
    fsconfig syscalls may occur in a different userns than that which
    called fsopen. That means that resolving the uid/gid via
    current_user_ns() can save a kuid that isn't mapped in the associated
    namespace when the filesystem is finally mounted. This means that it
    is possible for an unprivileged user to create files owned by any
    group in a tmpfs mount (since we can set the SUID bit on the tmpfs
    directory), or a tmpfs that is owned by any user, including the root
    group/user."

    The contract for {g,u}id mount options and {g,u}id values in general set
    from userspace has always been that they are translated according to the
    caller's idmapping. In so far, tmpfs has been doing the correct thing.
    But since tmpfs is mountable in unprivileged contexts it is also
    necessary to verify that the resulting {k,g}uid is representable in the
    namespace of the superblock to avoid such bugs as above.

    The new mount api's cross-namespace delegation abilities are already
    widely used. After having talked to a bunch of userspace this is the
    most faithful solution with minimal regression risks. I know of one
    users - systemd - that makes use of the new mount api in this way and
    they don't set unresolable {g,u}ids. So the regression risk is minimal.

    Link: https://lore.kernel.org/lkml/CALxfFW4BXhEwxR0Q5LSkg-
8Vb4r2MONKCcUCVioehXQKr35eHg@mail.gmail.com
    Fixes: f32356261d44 ("vfs: Convert ramfs, shmem, tmpfs, devtmpfs, rootfs to use the new
mount API")
    Reviewed-by: "Seth Forshee (DigitalOcean)" <sforshee@kernel.org>
    Reported-by: Seth Jenkins <sethjenkins@google.com>
    Message-Id: <20230801-vfs-fs_context-uidgid-v1-1-daf46a050bbf@kernel.org>
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit dfebf24b3aa008cef3ab1b8d4b3317886cb98e95
Author: Matthew Wilcox (Oracle) <willy@infradead.org>
Date:   Fri Jun 2 18:09:11 2023 -0400

    iomap: Remove large folio handling in iomap_invalidate_folio()

    [ Upstream commit a221ab717c43147f728d93513923ba3528f861bf ]

    We do not need to release the iomap_page in iomap_invalidate_folio()
```

```
    to allow the folio to be split.  The splitting code will call
    ->release_folio() if there is still per-fs private data attached to
    the folio.  At that point, we will check if the folio is still dirty
    and decline to release the iomap_page.  It is possible to trigger the
    warning in perfectly legitimate circumstances (eg if a disk read fails,
    we do a partial write to the folio, then we truncate the folio), which
    will cause those writes to be lost.

    Fixes: 60d8231089f0 ("iomap: Support large folios in invalidatepage")
    Signed-off-by: Matthew Wilcox (Oracle) <willy@infradead.org>
    Reviewed-by: Darrick J. Wong <djwong@kernel.org>
    Reviewed-by: Christoph Hellwig <hch@lst.de>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e3f23db0ef967a12faeeb5c28f595079f69f17d0
Author: Wang Ming <machel@vivo.com>
Date:   Thu Jul 13 20:05:42 2023 +0800

    fs: Fix error checking for d_hash_and_lookup()

    [ Upstream commit 0d5a4f8f775ff990142cdc810a84eae078589d27 ]

    The d_hash_and_lookup() function returns error pointers or NULL.
    Most incorrect error checks were fixed, but the one in int path_pts()
    was forgotten.

    Fixes: eedf265aa003 ("devpts: Make each mount of devpts an independent filesystem.")
    Signed-off-by: Wang Ming <machel@vivo.com>
    Message-Id: <20230713120555.7025-1-machel@vivo.com>
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 050287ef088a7d8a252a3ee6c917dc841cd14c58
Author: Wen Yang <wenyang.linux@foxmail.com>
Date:   Sun Jul 9 14:54:51 2023 +0800

    eventfd: prevent underflow for eventfd semaphores

    [ Upstream commit 758b492047816a3158d027e9fca660bc5bcf20bf ]

    For eventfd with flag EFD_SEMAPHORE, when its ctx->count is 0, calling
    eventfd_ctx_do_read will cause ctx->count to overflow to ULLONG_MAX.

    An underflow can happen with EFD_SEMAPHORE eventfds in at least the
    following three subsystems:

    (1) virt/kvm/eventfd.c
    (2) drivers/vfio/virqfd.c
    (3) drivers/virt/acrn/irqfd.c

    where (2) and (3) are just modeled after (1). An eventfd must be
    specified for use with the KVM_IRQFD ioctl(). This can also be an
    EFD_SEMAPHORE eventfd. When the eventfd count is zero or has been
    decremented to zero an underflow can be triggered when the irqfd is shut
    down by raising the KVM_IRQFD_FLAG_DEASSIGN flag in the KVM_IRQFD
    ioctl():

            // ctx->count == 0
            kvm_vm_ioctl()
            -> kvm_irqfd()
               -> kvm_irqfd_deassign()
                  -> irqfd_deactivate()
                     -> irqfd_shutdown()
                        -> eventfd_ctx_remove_wait_queue(&cnt)
                           -> eventfd_ctx_do_read(&cnt)
```

```
    Userspace polling on the eventfd wouldn't notice the underflow because 1
    is always returned as the value from eventfd_read() while ctx->count
    would've underflowed. It's not a huge deal because this should only be
    happening when the irqfd is shutdown but we should still fix it and
    avoid the spurious wakeup.

    Fixes: cb289d6244a3 ("eventfd - allow atomic read and waitqueue remove")
    Signed-off-by: Wen Yang <wenyang.linux@foxmail.com>
    Cc: Alexander Viro <viro@zeniv.linux.org.uk>
    Cc: Jens Axboe <axboe@kernel.dk>
    Cc: Christian Brauner <brauner@kernel.org>
    Cc: Christoph Hellwig <hch@lst.de>
    Cc: Dylan Yudaken <dylany@fb.com>
    Cc: David Woodhouse <dwmw@amazon.co.uk>
    Cc: Matthew Wilcox <willy@infradead.org>
    Cc: linux-fsdevel@vger.kernel.org
    Cc: linux-kernel@vger.kernel.org
    Message-Id: <tencent_7588DFD1F365950A757310D764517A14B306@qq.com>
    [brauner: rewrite commit message and add explanation how this underflow can happen]
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit 6ba0372b101504a44a9927c719a0698018140714
Author: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
Date:   Mon Jul 3 16:42:21 2023 +0200

    splice: fsnotify_access(in), fsnotify_modify(out) on success in tee

    [ Upstream commit 576d498e0ac5caff2d9f6312573ab54d98f12d32 ]

    Same logic applies here: this can fill up the pipe, and pollers that rely
    on getting IN_MODIFY notifications never wake up.

    Fixes: 983652c69199 ("splice: report related fsnotify events")
    Link: https://lore.kernel.org/linux-
fsdevel/jbyihkyk5dtaohdwjyivambb2gffyjs3dodpofafnkkunxq7bu@jngkdxx65pux/t/#u
    Link: https://bugs.debian.org/1039488
    Signed-off-by: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
    Acked-by: Jan Kara <jack@suse.cz>
    Reviewed-by: Amir Goldstein <amir73il@gmail.com>
    Message-Id:
<10d76dd8c85017ae3cd047c9b9a32e26daefdaa2.1688393619.git.nabijaczleweli@nabijaczleweli.xyz>
    Signed-off-by: Christian Brauner <brauner@kernel.org>
    Signed-off-by: Sasha Levin <sashal@kernel.org>

commit e2f5ea718f0b05446f369854b0e5cf9dc3fc9e9a
Author: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
Date:   Mon Jul 3 16:42:17 2023 +0200

    splice: fsnotify_access(fd)/fsnotify_modify(fd) in vmsplice

    [ Upstream commit 7f0f1ea069e52d5a16921abd59377a7da6c25149 ]

    Same logic applies here: this can fill up the pipe and pollers that rely
    on getting IN_MODIFY notifications never wake up.

    Fixes: 983652c69199 ("splice: report related fsnotify events")
    Link: https://lore.kernel.org/linux-
fsdevel/jbyihkyk5dtaohdwjyivambb2gffyjs3dodpofafnkkunxq7bu@jngkdxx65pux/t/#u
    Link: https://bugs.debian.org/1039488
    Signed-off-by: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
    Acked-by: Jan Kara <jack@suse.cz>
    Reviewed-by: Amir Goldstein <amir73il@gmail.com>
    Message-Id:
```

```
 <8d9ad5acb9c5c1dd2376a2ff5da6ac3183115389.1688393619.git.nabijaczleweli@nabijaczleweli.xyz>
     Signed-off-by: Christian Brauner <brauner@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 5afd3f424aa9a28c7b6862ce8c1482c03570deeb
 Author: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
 Date:   Mon Jul 3 16:42:13 2023 +0200

     splice: always fsnotify_access(in), fsnotify_modify(out) on success

     [ Upstream commit 12ee4b66af34f8e72f3b2fd93a946a955efe7c86 ]

     The current behaviour caused an asymmetry where some write APIs
     (write, sendfile) would notify the written-to/read-from objects,
     but splice wouldn't.

     This affected userspace which uses inotify, most notably coreutils
     tail -f, to monitor pipes.
     If the pipe buffer had been filled by a splice-family function:
       * tail wouldn't know and thus wouldn't service the pipe, and
       * all writes to the pipe would block because it's full,
     thus service was denied.
     (For the particular case of tail -f this could be worked around
      with ---disable-inotify.)

     Fixes: 983652c69199 ("splice: report related fsnotify events")
     Link: https://lore.kernel.org/linux-
 fsdevel/jbyihkyk5dtaohdwjyivambb2gffyjs3dodpofafnkkunxq7bu@jngkdxx65pux/t/#u
     Link: https://bugs.debian.org/1039488
     Signed-off-by: Ahelenia Ziemiańska <nabijaczleweli@nabijaczleweli.xyz>
     Acked-by: Jan Kara <jack@suse.cz>
     Reviewed-by: Amir Goldstein <amir73il@gmail.com>
     Message-Id:
 <604ec704d933e0e0121d9e107ce914512e045fad.1688393619.git.nabijaczleweli@nabijaczleweli.xyz>
     Signed-off-by: Christian Brauner <brauner@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit fcc2b8ac9e49b55d75d70879bd29a013b5c4d049
 Author: Matthew Wilcox <willy@infradead.org>
 Date:   Sun Jun 4 12:16:06 2023 +0100

     reiserfs: Check the return value from __getblk()

     [ Upstream commit ba38980add7ffc9e674ada5b4ded4e7d14e76581 ]

     __getblk() can return a NULL pointer if we run out of memory or if we
     try to access beyond the end of the device; check it and handle it
     appropriately.

     Signed-off-by: Matthew Wilcox (Oracle) <willy@infradead.org>
     Link:
 https://lore.kernel.org/lkml/CAFcO6XOacq3hscbXevPQP7sXRoYFz34ZdKPYjmd6k5sZuhGFDw@mail.gmail.
 com/
     Tested-by: butt3rflyh4ck <butterflyhuangxx@gmail.com>
     Fixes: 1da177e4c3f4 ("Linux-2.6.12-rc2") # probably introduced in 2002
     Acked-by: Edward Shishkin <edward.shishkin@gmail.com>
     Signed-off-by: Christian Brauner <brauner@kernel.org>
     Signed-off-by: Sasha Levin <sashal@kernel.org>

 commit 61e2589decb79b23c253c30b22355426ebca9ff1
 Author: Sabrina Dubroca <sd@queasysnail.net>
 Date:   Mon Sep 4 10:56:04 2023 +0200

     Revert "net: macsec: preserve ingress frame ordering"
```

```
    commit d3287e4038ca4f81e02067ab72d087af7224c68b upstream.

    This reverts commit ab046a5d4be4c90a3952a0eae75617b49c0cb01b.

    It was trying to work around an issue at the crypto layer by excluding
    ASYNC implementations of gcm(aes), because a bug in the AESNI version
    caused reordering when some requests bypassed the cryptd queue while
    older requests were still pending on the queue.

    This was fixed by commit 38b2f68b4264 ("crypto: aesni - Fix cryptd
    reordering problem on gcm"), which pre-dates ab046a5d4be4.

    Herbert Xu confirmed that all ASYNC implementations are expected to
    maintain the ordering of completions wrt requests, so we can use them
    in MACsec.

    On my test machine, this restores the performance of a single netperf
    instance, from 1.4Gbps to 4.4Gbps.

    Link:
https://lore.kernel.org/netdev/9328d206c5d9f9239cae27e62e74de40b258471d.1692279161.git.sd@qu
easysnail.net/T/
    Link: https://lore.kernel.org/netdev/1b0cec71-d084-8153-2ba4-72ce71abeb65@byu.edu/
    Link: https://lore.kernel.org/netdev/d335ddaa-18dc-f9f0-17ee-9783d3b2ca29@mailbox.tu-
dresden.de/
    Fixes: ab046a5d4be4 ("net: macsec: preserve ingress frame ordering")
    Signed-off-by: Sabrina Dubroca <sd@queasysnail.net>
    Link:
https://lore.kernel.org/r/11c952469d114db6fb29242e1d9545e61f52f512.1693757159.git.sd@queasys
nail.net
    Signed-off-by: Paolo Abeni <pabeni@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f3de44835ce00bb06e03e41e7e088b927f6993ea
Author: Vidya Sagar <vidyas@nvidia.com>
Date:   Mon Jun 19 15:56:04 2023 +0530

    Revert "PCI: tegra194: Enable support for 256 Byte payload"

    commit ebfde1584d9f037b6309fc682c96e22dac7bcb7a upstream.

    After commit 4fb8e46c1bc4 ("PCI: tegra194: Enable support for 256 Byte
    payload"), we initialize MPS=256 for tegra194 Root Ports before enumerating
    the hierarchy.

    Consider an Endpoint that supports only MPS=128.  In the default situation
    (CONFIG_PCIE_BUS_DEFAULT set and no "pci=pcie_bus_*" parameter), Linux
    tries to configure the MPS of every device to match the upstream bridge.
    If the Endpoint is directly below the Root Port, Linux can reduce the Root
    Port MPS to 128 to match the Endpoint.  But if there's a switch in the
    middle, Linux doesn't reduce the Root Port MPS because other devices below
    the switch may already be configured with MPS larger than 128.

    This scenario results in uncorrectable Malformed TLP errors if the Root
    Port sends TLPs with payloads larger than 128 bytes.  These errors can
    be avoided by using the "pci=pcie_bus_safe" parameter, but it doesn't
    seem to be a good idea to always have this parameter even for basic
    functionality to work.

    Revert commit 4fb8e46c1bc4 ("PCI: tegra194: Enable support for 256 Byte
    payload") so the Root Ports default to MPS=128, which all devices
    support.

    If peer-to-peer DMA is not required, one can use "pci=pcie_bus_perf" to
    get the benefit of larger MPS settings.
```

```
    [bhelgaas: commit log; kwilczynski: retain "u16 val_16" declaration at
    the top, add missing acked by tag]
    Fixes: 4fb8e46c1bc4 ("PCI: tegra194: Enable support for 256 Byte payload")
    Link: https://lore.kernel.org/linux-pci/20230619102604.3735001-1-vidyas@nvidia.com
    Signed-off-by: Vidya Sagar <vidyas@nvidia.com>
    Signed-off-by: Krzysztof Wilczyński <kwilczynski@kernel.org>
    Acked-by: Jon Hunter <jonathanh@nvidia.com>
    Cc: stable@vger.kernel.org # v6.0-rc1+
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 75ca73b41647d4a311d78583abc72097aa08a864
Author: Jaegeuk Kim <jaegeuk@kernel.org>
Date:   Fri Aug 4 12:15:34 2023 -0700

    Revert "f2fs: clean up w/ sbi->log_sectors_per_block"

    commit 579c7e41507e85dc3eedf998a3dca14a2a1526ad upstream.

    This reverts commit bfd476623999118d9c509cb0fa9380f2912bc225.

    Shinichiro Kawasaki reported:

    When I ran workloads on f2fs using v6.5-rcX with fixes [1][2] and a zoned block
    devices with 4kb logical block size, I observe mount failure as follows. When
    I revert this commit, the failure goes away.

    [  167.781975][ T1555] F2FS-fs (dm-0): IO Block Size:        4 KB
    [  167.890728][ T1555] F2FS-fs (dm-0): Found nat_bits in checkpoint
    [  171.482588][ T1555] F2FS-fs (dm-0): Zone without valid block has non-zero write
pointer. Reset the write pointer: wp[0x1300,0x8]
    [  171.496000][ T1555] F2FS-fs (dm-0): (0) : Unaligned zone reset attempted (block
280000 + 80000)
    [  171.505037][ T1555] F2FS-fs (dm-0): Discard zone failed:  (errno=-5)

    The patch replaced "sbi->log_blocksize - SECTOR_SHIFT" with
    "sbi->log_sectors_per_block". However, I think these two are not equal when the
    device has 4k logical block size. The former uses Linux kernel sector size 512
    byte. The latter use 512b sector size or 4kb sector size depending on the
    device. mkfs.f2fs obtains logical block size via BLKSSZGET ioctl from the device
    and reflects it to the value sbi->log_sector_size_per_block. This causes
    unexpected write pointer calculations in check_zone_write_pointer(). This
    resulted in unexpected zone reset and the mount failure.

    [1] https://lkml.kernel.org/linux-f2fs-devel/20230711050101.GA19128@lst.de/
    [2] https://lore.kernel.org/linux-f2fs-devel/20230804091556.2372567-1-
shinichiro.kawasaki@wdc.com/

    Cc: stable@vger.kernel.org
    Reported-by: Shinichiro Kawasaki <shinichiro.kawasaki@wdc.com>
    Fixes: bfd476623999 ("f2fs: clean up w/ sbi->log_sectors_per_block")
    Reviewed-by: Chao Yu <chao@kernel.org>
    Signed-off-by: Jaegeuk Kim <jaegeuk@kernel.org>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit ec3c547fba2d3dde1f0a7c622ed3fb2cae0f4bc6
Author: Miklos Szeredi <mszeredi@redhat.com>
Date:   Mon Aug 14 13:05:30 2023 +0200

    Revert "fuse: in fuse_flush only wait if someone wants the return code"

    commit 91ec6c85599b60c00caf4e9a9d6c4d6e5dd5e93c upstream.

    This reverts commit 5a8bee63b10f6f2f52f6d22e109a4a147409842a.
```

```
    Jürg Billeter reports the following regression:

      Since v6.3-rc1 commit 5a8bee63b1 ("fuse: in fuse_flush only wait if
      someone wants the return code") `fput()` is called asynchronously if a
      file is closed as part of a process exiting, i.e., if there was no
      explicit `close()` before exit.

      If the file was open for writing, also `put_write_access()` is called
      asynchronously as part of the async `fput()`.

      If that newly written file is an executable, attempting to `execve()` the
      new file can fail with `ETXTBSY` if it's called after the writer process
      exited but before the async `fput()` has run.

    Reported-and-tested-by: "Jürg Billeter" <j@bitron.ch>
    Cc: <stable@vger.kernel.org> # v6.3
    Link:
https://lore.kernel.org/all/4f66cded234462964899f2a661750d6798a57ec0.camel@bitron.ch/
    Signed-off-by: Miklos Szeredi <mszeredi@redhat.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit f48c2c758e60807d508cd9271b292208273f0d05
Author: Werner Sembach <wse@tuxedocomputers.com>
Date:   Wed Jul 12 11:56:51 2023 -0700

    Input: i8042 - add quirk for TUXEDO Gemini 17 Gen1/Clevo PD70PN

    commit eb09074bdb05ffd6bfe77f8b4a41b76ef78c997b upstream.

    The touchpad of this device is both connected via PS/2 and i2c. This causes
    strange behavior when both driver fight for control. The easy fix is to
    prevent the PS/2 driver from accessing the mouse port as the full feature
    set of the touchpad is only supported in the i2c interface anyway.

    The strange behavior in this case is, that when an external screen is
    connected and the notebook is closed, the pointer on the external screen is
    moving to the lower right corner. When the notebook is opened again, this
    movement stops, but the touchpad clicks are unresponsive afterwards until
    reboot.

    Signed-off-by: Werner Sembach <wse@tuxedocomputers.com>
    Cc: stable@vger.kernel.org
    Link: https://lore.kernel.org/r/20230607173331.851192-1-wse@tuxedocomputers.com
    Signed-off-by: Dmitry Torokhov <dmitry.torokhov@gmail.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit 98084525a959510ff59f4c32aba743e1d128d7ce
Author: Max Chou <max.chou@realtek.com>
Date:   Mon Aug 7 19:42:59 2023 +0800

    Bluetooth: btrtl: Load FW v2 otherwise FW v1 for RTL8852C

    commit bd003fb338afee97c76f13c3e9144a7e4ad37179 upstream.

    In this commit, prefer to load FW v2 if available. Fallback to FW v1
    otherwise. This behavior is only for RTL8852C.

    Fixes: 9a24ce5e29b1 ("Bluetooth: btrtl: Firmware format v2 support")
    Cc: stable@vger.kernel.org
    Suggested-by: Juerg Haefliger <juerg.haefliger@canonical.com>
    Tested-by: Hilda Wu <hildawu@realtek.com>
    Signed-off-by: Max Chou <max.chou@realtek.com>
    Signed-off-by: Luiz Augusto von Dentz <luiz.von.dentz@intel.com>
    [juergh: Adjusted context due to missing .hw_info struct element]
    Signed-off-by: Juerg Haefliger <juerg.haefliger@canonical.com>
```

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

commit 46c46557af238467b6e82c1e70d21ca9013d5223
Author: Jordan Rife <jrife@google.com>
Date:   Mon Aug 21 16:45:23 2023 -0500

    net: Avoid address overwrite in kernel_connect

    commit 0bdf399342c5acbd817c9098b6c7ed21f1974312 upstream.

    BPF programs that run on connect can rewrite the connect address. For
    the connect system call this isn't a problem, because a copy of the address
    is made when it is moved into kernel space. However, kernel_connect
    simply passes through the address it is given, so the caller may observe
    its address value unexpectedly change.

    A practical example where this is problematic is where NFS is combined
    with a system such as Cilium which implements BPF-based load balancing.
    A common pattern in software-defined storage systems is to have an NFS
    mount that connects to a persistent virtual IP which in turn maps to an
    ephemeral server IP. This is usually done to achieve high availability:
    if your server goes down you can quickly spin up a replacement and remap
    the virtual IP to that endpoint. With BPF-based load balancing, mounts
    will forget the virtual IP address when the address rewrite occurs
    because a pointer to the only copy of that address is passed down the
    stack. Server failover then breaks, because clients have forgotten the
    virtual IP address. Reconnects fail and mounts remain broken. This patch
    was tested by setting up a scenario like this and ensuring that NFS
    reconnects worked after applying the patch.

    Signed-off-by: Jordan Rife <jrife@google.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bb3d1e0b966331fc25a939c52982b464cca16f7a
Author: Jarkko Sakkinen <jarkko@kernel.org>
Date:   Mon Sep 4 21:12:10 2023 +0300

    tpm: Enable hwrng only for Pluton on AMD CPUs

    commit 8f7f35e5aa6f2182eabcfa3abef4d898a48e9aa8 upstream.

    The vendor check introduced by commit 554b841d4703 ("tpm: Disable RNG for
    all AMD fTPMs") doesn't work properly on a number of Intel fTPMs.  On the
    reported systems the TPM doesn't reply at bootup and returns back the
    command code. This makes the TPM fail probe on Lenovo Legion Y540 laptop.

    Since only Microsoft Pluton is the only known combination of AMD CPU and
    fTPM from other vendor, disable hwrng otherwise. In order to make sysadmin
    aware of this, print also info message to the klog.

    Cc: stable@vger.kernel.org
    Fixes: 554b841d4703 ("tpm: Disable RNG for all AMD fTPMs")
    Reported-by: Todd Brandt <todd.e.brandt@intel.com>
    Closes: https://bugzilla.kernel.org/show_bug.cgi?id=217804
    Reported-by: Patrick Steinhardt <ps@pks.im>
    Reported-by: Raymond Jay Golo <rjgolo@gmail.com>
    Reported-by: Ronan Pigott <ronan@rjp.ie>
    Reviewed-by: Jerry Snitselaar <jsnitsel@redhat.com>
    Signed-off-by: Jarkko Sakkinen <jarkko@kernel.org>
    Cc: Thorsten Leemhuis <regressions@leemhuis.info>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit b9dc1ea031439e095382c63c80d923f72052b78b
Author: Doug Smythies <dsmythies@telus.net>
```

```
Date:   Sun Aug 20 13:46:49 2023 -0700

    cpufreq: intel_pstate: set stale CPU frequency to minimum

    commit d51847acb018d83186e4af67bc93f9a00a8644f7 upstream.

    The intel_pstate CPU frequency scaling driver does not
    use policy->cur and it is 0.
    When the CPU frequency is outdated arch_freq_get_on_cpu()
    will default to the nominal clock frequency when its call to
    cpufreq_quick_getpolicy_cur returns the never updated 0.
    Thus, the listed frequency might be outside of currently
    set limits. Some users are complaining about the high
    reported frequency, albeit stale, when their system is
    idle and/or it is above the reduced maximum they have set.

    This patch will maintain policy_cur for the intel_pstate
    driver at the current minimum CPU frequency.

    Reported-by: Yang Jie <yang.jie@linux.intel.com>
    Closes: https://bugzilla.kernel.org/show_bug.cgi?id=217597
    Signed-off-by: Doug Smythies <dsmythies@telus.net>
    [ rjw: White space damage fixes and comment adjustment ]
    Signed-off-by: Rafael J. Wysocki <rafael.j.wysocki@intel.com>
    Signed-off-by: Keyon Jie <yang.jie@linux.intel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>

commit bf2681249928fbd44975af51d4798f29f7bc6bfe
Author: Hamza Mahfooz <hamza.mahfooz@amd.com>
Date:   Fri Aug 4 11:13:04 2023 -0400

    drm/amd/display: ensure async flips are only accepted for fast updates

    commit a7c0cad0dc060bb77e9c9d235d68441b0fc69507 upstream.

    We should be checking to see if async flips are supported in
    amdgpu_dm_atomic_check() (i.e. not dm_crtc_helper_atomic_check()). Also,
    async flipping isn't supported if a plane's framebuffer changes memory
    domains during an atomic commit. So, move the check from
    dm_crtc_helper_atomic_check() to amdgpu_dm_atomic_check() and check if
    the memory domain has changed in amdgpu_dm_atomic_check().

    Cc: stable@vger.kernel.org
    Link: https://gitlab.freedesktop.org/drm/amd/-/issues/2733
    Fixes: c1e18c44dc7f ("drm/amd/display: only accept async flips for fast updates")
    Reviewed-by: Harry Wentland <harry.wentland@amd.com>
    Signed-off-by: Hamza Mahfooz <hamza.mahfooz@amd.com>
    Signed-off-by: Alex Deucher <alexander.deucher@amd.com>
    Reported-by: Michael Larabel <Michael@MichaelLarabel.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```