Fixing Coverity Bugs in the Linux Kernel

Gustavo A. R. Silva
gustavo@embeddedor.com

Supported by The Linux Foundation's Core Infrastructure Initiative

Kernel Recipes 2017

Overview

- What am I doing to get involved?
- How am I doing it?
- Response from the community.
- The future of this project.

Coverity

- Static code analyzer.
- Performs analysis without running the code.

Interface

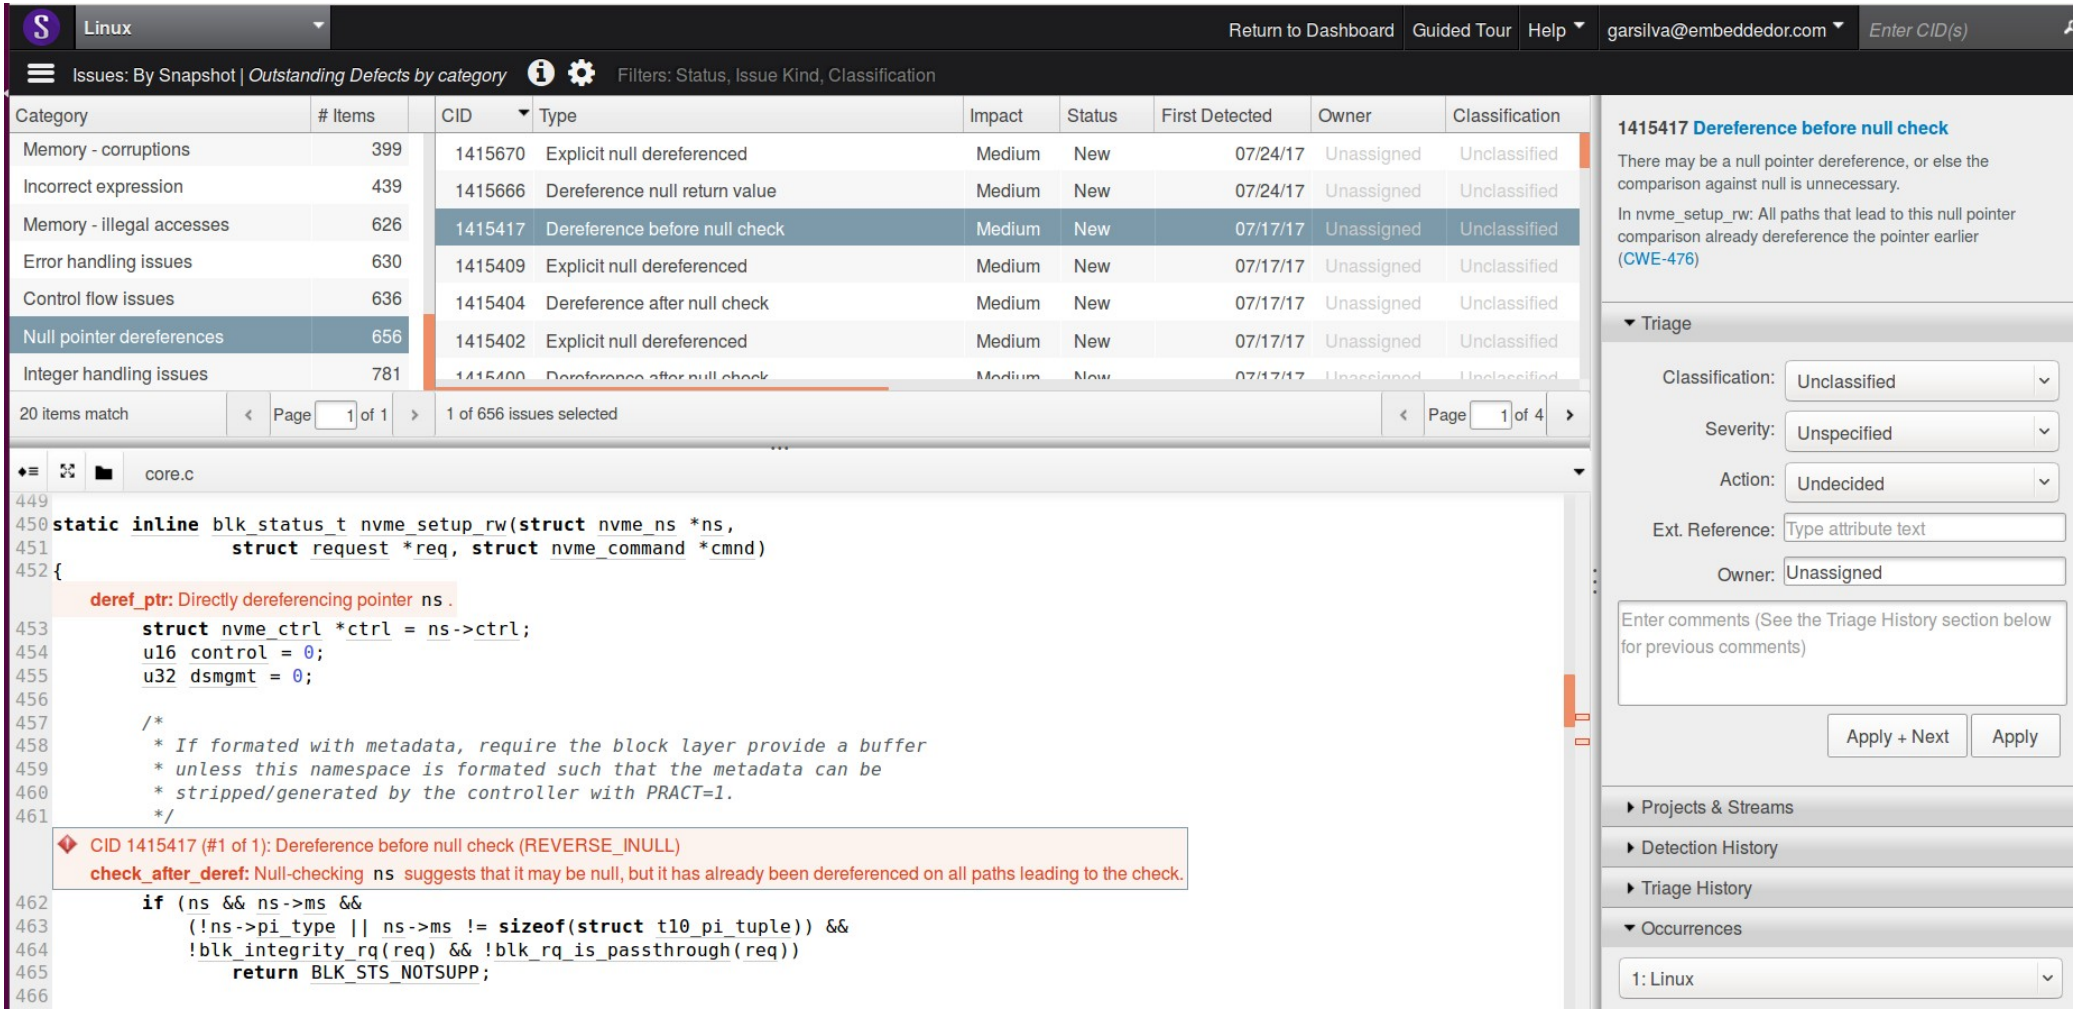
Linux
Return to Dashboard
Guided Tour
Help
garsilva@embeddedor.com
Enter CID(s)
Issues: By Snapshot | Outstanding Defects by category
Filters: Status, Issue Kind, Classification
Category
Items
Memory - corruptions 399
Incorrect expression 439
Memory - illegal accesses 626
Error handling issues 630
Control flow issues 636
Null pointer dereferences 656
Integer handling issues 781
CID
Type
Impact
Status
First Detected
Owner
Classification
1415670 Explicit null dereferenced Medium New 07/24/17 Unassigned Unclassified
1415666 Dereference null return value Medium New 07/24/17 Unassigned Unclassified
1415417 Dereference before null check Medium New 07/17/17 Unassigned Unclassified
1415409 Explicit null dereferenced Medium New 07/17/17 Unassigned Unclassified
1415404 Dereference after null check Medium New 07/17/17 Unassigned Unclassified
1415402 Explicit null dereferenced Medium New 07/17/17 Unassigned Unclassified
20 items match
Page 1 of 1
1 of 656 issues selected
Page 1 of 4
core.c
449
450 static inline blk_status_t nvme_setup_rw(struct nvme_ns *ns,
451 struct request *req, struct nvme_command *cmnd)
452 {
deref_ptr: Directly dereferencing pointer ns .
453 struct nvme_ctrl *ctrl = ns->ctrl;
454 u16 control = 0;
455 u32 dsmgmt = 0;
456
457 /*
458 * If formated with metadata, require the block layer provide a buffer
459 * unless this namespace is formated such that the metadata can be
460 * stripped/generated by the controller with PRACT=1.
461 */
CID 1415417 (#1 of 1): Dereference before null check (REVERSE_INULL)
check_after_deref: Null-checking ns suggests that it may be null, but it has already been dereferenced on all paths leading to the check.
462 if (ns && ns->ms &&
463 (!ns->pi_type || ns->ms != sizeof(struct t10_pi_tuple)) &&
464 !blk_integrity_rq(req) && !blk_rq_is_passthrough(req))
465 return BLK_STS_NOTSUPP;
466
1415417 Dereference before null check
There may be a null pointer dereference, or else the comparison against null is unnecessary.
In nvme_setup_rw: All paths that lead to this null pointer comparison already dereference the pointer earlier (CWE-476)
Triage
Classification: Unclassified
Severity: Unspecified
Action: Undecided
Ext. Reference: Type attribute text
Owner: Unassigned
Enter comments (See the Triage History section below for previous comments)
Apply + Next
Apply
Projects & Streams
Detection History
Triage History
Occurrences
1: Linux

Last Coverity report

High impact issues

- 216 illegal memory accesses.

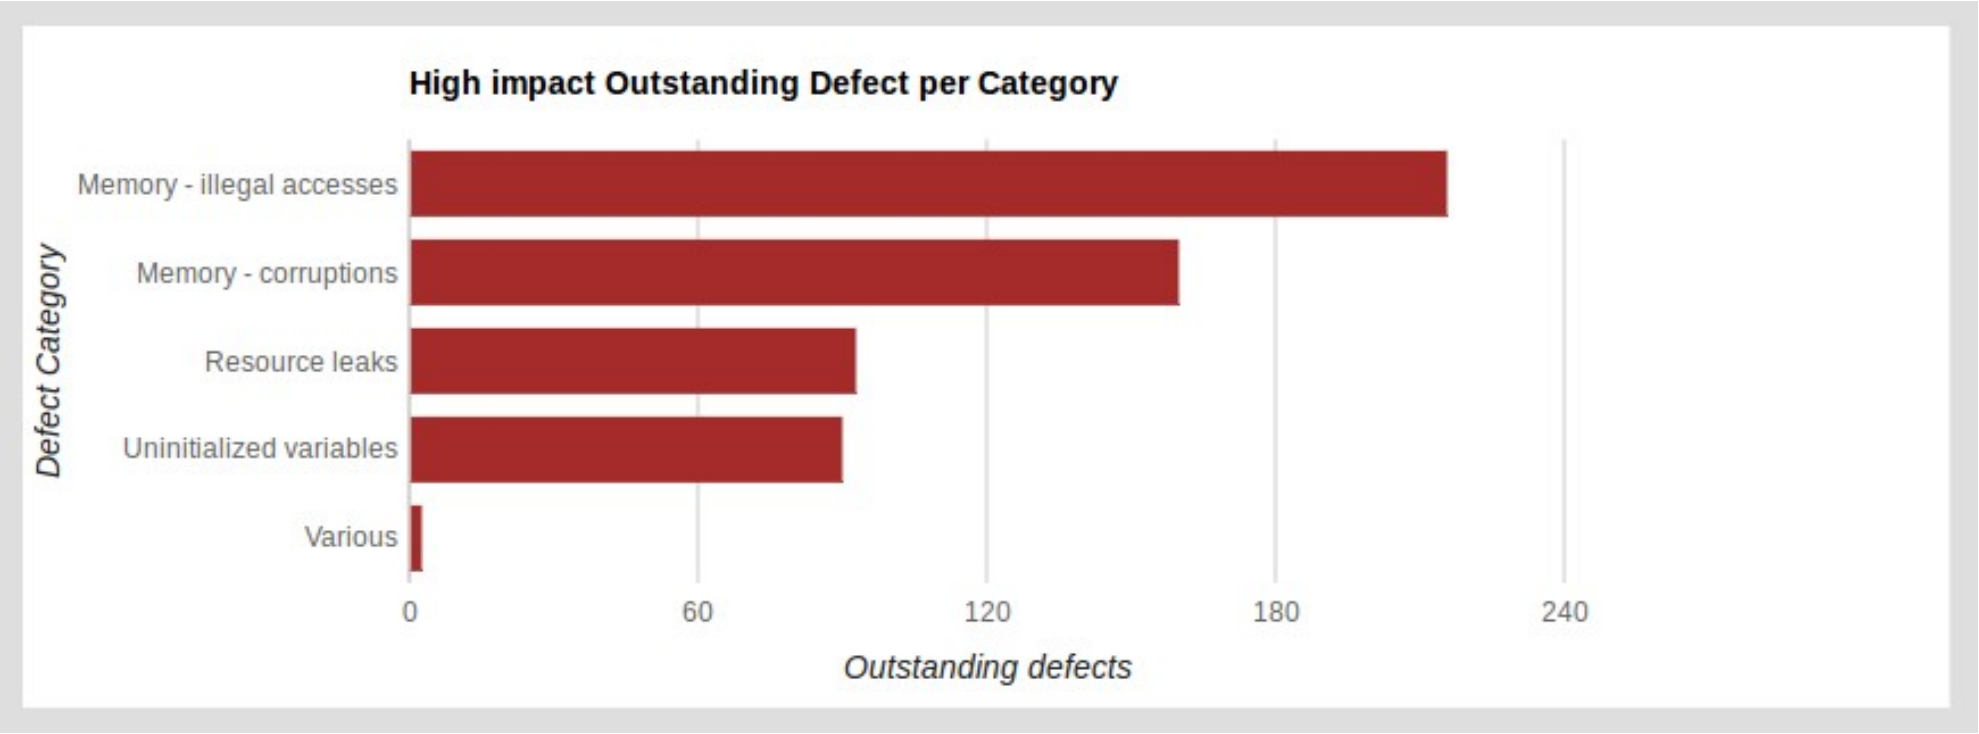

Illegal memory accesses

- Out-of-bounds access.
- Use after free.
- Buffer not null terminated.
- Negative array index read.

Medium impact issues

- 310 null pointer dereferences.

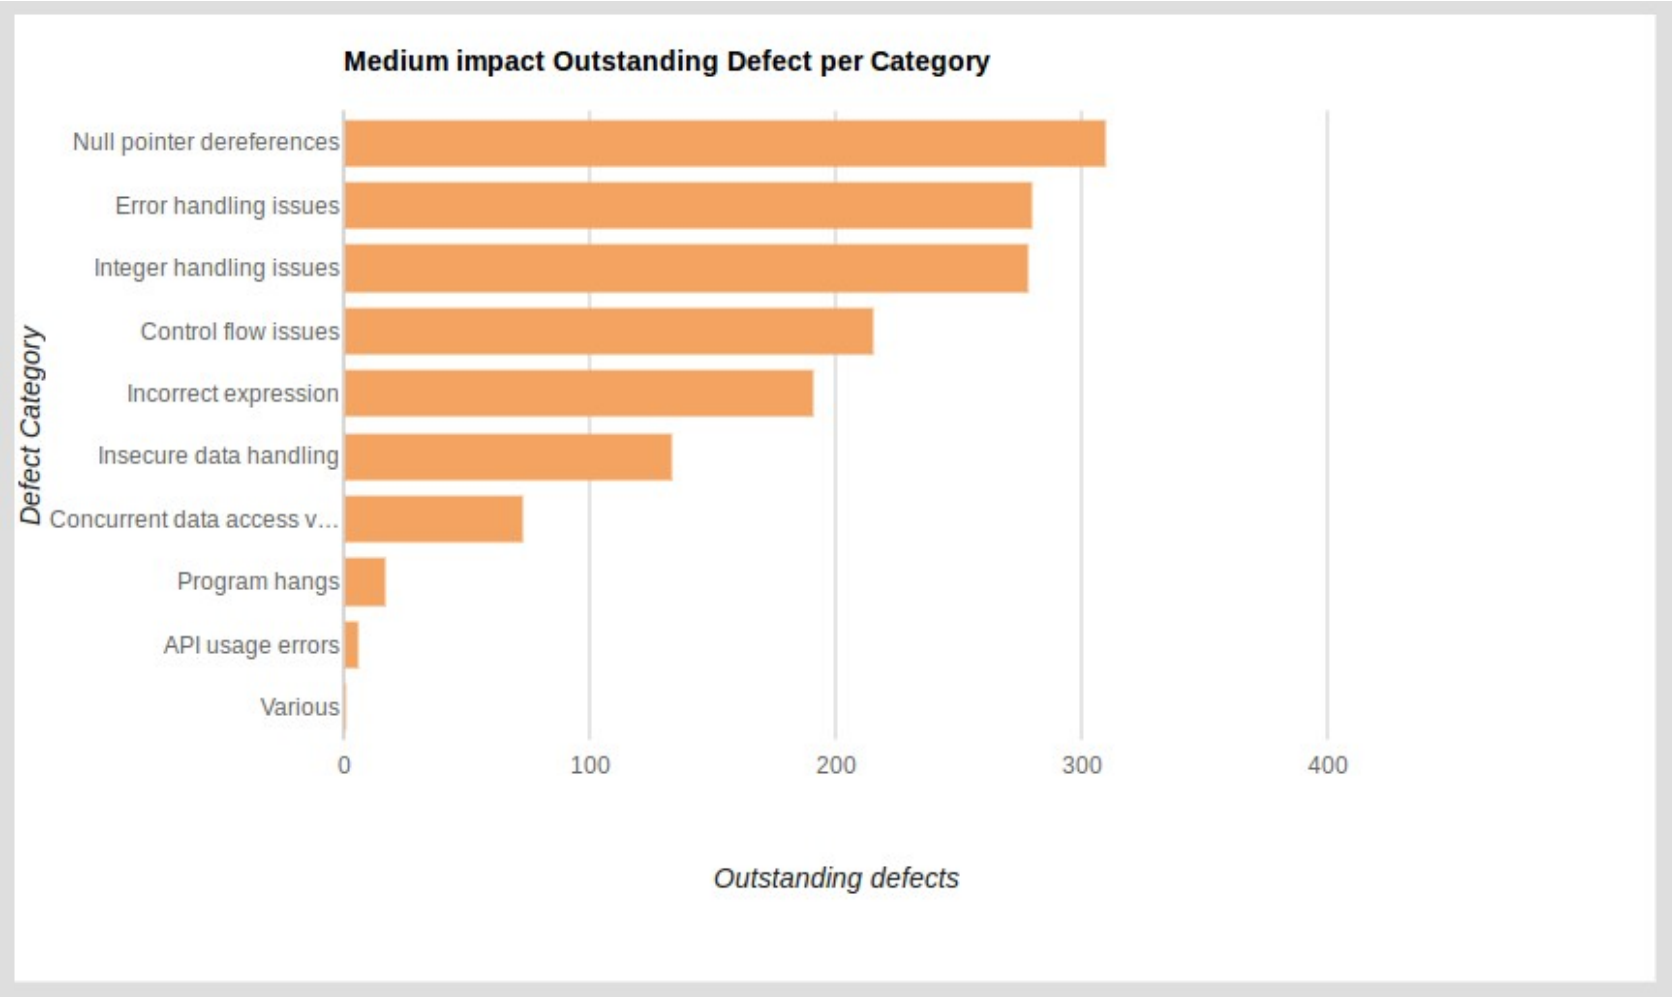

Some examples

From ‘Missing break in switch’ to Code refactoring

- drivers/usb/misc/usbtest.c

```
/* take the first altsetting with in-bulk + out-bulk;
 * ignore other endpoints and altsettings.
 */
for (ep = 0; ep < alt->desc.bNumEndpoints; ep++) {
        struct usb_host_endpoint        *e;

        e = alt->endpoint + ep;
        switch (usb_endpoint_type(&e->desc)) {
        case USB_ENDPOINT_XFER_BULK:
                break;
        case USB_ENDPOINT_XFER_INT:
                if (dev->info->intr)
                        goto try_intr;
        case USB_ENDPOINT_XFER_ISOC:
                if (dev->info->iso)
                        goto try_iso;
                /* FALLTHROUGH */
        default:
                continue;
        }
        if (usb_endpoint_dir_in(&e->desc)) {
                if (!in)
                        in = e;
        } else {
                if (!out)
                        out = e;
        }
        continue;
```

From ‘missing break in switch’ to code refactoring

- It turned out to be a missing continue.

```
diff --git a/drivers/usb/misc/usbtest.c b/drivers/usb/misc/usbtest.c
index 17c0810..26ae5d1 100644
--- a/drivers/usb/misc/usbtest.c
+++ b/drivers/usb/misc/usbtest.c
@@ -159,6 +159,7 @@ get_endpoints(struct usbtest_dev *dev, struct usb_interface *intf)
                        case USB_ENDPOINT_XFER_INT:
                                if (dev->info->intr)
                                        goto try_intr;
+                               continue;
                        case USB_ENDPOINT_XFER_ISOC:
                                if (dev->info->iso)
                                        goto try_iso;
```

There was some room for improvement.

```
diff --git a/drivers/usb/misc/usbtest.c b/drivers/usb/misc/usbtest.c
index 26ae5d1..eee82ca 100644
--- a/drivers/usb/misc/usbtest.c
+++ b/drivers/usb/misc/usbtest.c
@@ -124,6 +124,20 @@ static struct usb_device *testdev_to_usbdev(struct usbtest_dev *test

 /*-------------------------------------------------------------------------*/

+static inline void endpoint_update(int edi,
+                                   struct usb_host_endpoint **in,
+                                   struct usb_host_endpoint **out,
+                                   struct usb_host_endpoint *e)
+{
+       if (edi) {
+               if (!*in)
+                       *in = e;
+       } else {
+               if (!*out)
+                       *out = e;
+       }
+}
+
 static int
 get_endpoints(struct usbtest_dev *dev, struct usb_interface *intf)
 {
@@ -151,47 +165,26 @@ get_endpoints(struct usbtest_dev *dev, struct usb_interface *intf)
                 */
```

Arguments in wrong order

- drivers/scsi/qedf/qedf_els.c

```
__fc_fill_fc_hdr(fc_hdr, FC_RCTL_ELS_REQ, sid, did,
                 FC_TYPE_ELS, FC_FC_FIRST_SEQ | FC_FC_END_SEQ |
                 FC_FC_SEQ_INIT, 0);
```

- include/scsi/fc_encode.h

```
static inline void __fc_fill_fc_hdr(struct fc_frame_header *fh,
                                    enum fc_rctl r_ctl,
                                    u32 did, u32 sid, enum fc_fh_type type,
                                    u32 f_ctl, u32 parm_offset)
```

The fix

```
diff --git a/drivers/scsi/qedf/qedf_els.c b/drivers/scsi/qedf/qedf_els.c
index c505d41..9062703 100644
--- a/drivers/scsi/qedf/qedf_els.c
+++ b/drivers/scsi/qedf/qedf_els.c
@@ -109,7 +109,7 @@ retry_els:
        did = fcport->rdata->ids.port_id;
        sid = fcport->sid;

-       __fc_fill_fc_hdr(fc_hdr, FC_RCTL_ELS_REQ, sid, did,
+       __fc_fill_fc_hdr(fc_hdr, FC_RCTL_ELS_REQ, did, sid,
                         FC_TYPE_ELS, FC_FC_FIRST_SEQ | FC_FC_END_SEQ |
                         FC_FC_SEQ_INIT, 0);
```

'Uninitialized scalar variable' turned out to be a copy/paste error

- drivers/scsi/libfc/fc_rport.c

```
fp = fc_frame_alloc(lport, sizeof(*rtv));
if (!fp) {
        rjt_data.reason = ELS_RJT_UNAB;
        rjt_data.reason = ELS_EXPL_INSUF_RES;
        fc_seq_els_rsp_send(in_fp, ELS_LS_RJT, &rjt_data);
        goto drop;
}
```

- include/scsi/libfc.h

```
struct fc_seq_els_data {
        enum fc_els_rjt_reason reason;
        enum fc_els_rjt_explan explan;
};
```

els_data->explan

```
void fc_seq_els_rsp_send(struct fc_frame *fp, enum fc_els_cmd els_cmd,
                         struct fc_seq_els_data *els_data)
{
        switch (els_cmd) {
        case ELS_LS_RJT:
                fc_seq_ls_rjt(fp, els_data->reason, els_data->explan);
                break;
        case ELS_LS_ACC:
                fc_seq_ls_acc(fp);
                break;
        case ELS_RRQ:
                fc_exch_els_rrq(fp);
                break;
        case ELS_REC:
                fc_exch_els_rec(fp);
                break;
        default:
                FC_LPORT_DBG(fr_dev(fp), "Invalid ELS CMD:%x\n", els_cmd);
        }
}
```

ELS_EXPL_INSUF_RES

Defined in 1 files:

include/uapi/scsi/fc/fc_els.h, line 212 *(as a enumerator)*

Referenced in 2 files:

drivers/scsi/libfc/fc_rport.c

- line 1425
- line 1650
- line 1866
- line 2000
- line 2112

include/uapi/scsi/fc/fc_els.h, line 212

Same pattern in all cases

```
rjt_data.reason = ELS_RJT_UNAB;
rjt_data.explan = ELS_EXPL_INSUF_RES;
```

The fix

```
diff --git a/drivers/scsi/libfc/fc_rport.c b/drivers/scsi/libfc/fc_rport.c
index b44c313..5203258 100644
--- a/drivers/scsi/libfc/fc_rport.c
+++ b/drivers/scsi/libfc/fc_rport.c
@@ -1422,7 +1422,7 @@ static void fc_rport_recv_rtv_req(struct fc_rport_priv *rdata,
        fp = fc_frame_alloc(lport, sizeof(*rtv));
        if (!fp) {
                rjt_data.reason = ELS_RJT_UNAB;
-               rjt_data.reason = ELS_EXPL_INSUF_RES;
+               rjt_data.explan = ELS_EXPL_INSUF_RES;
                fc_seq_els_rsp_send(in_fp, ELS_LS_RJT, &rjt_data);
                goto drop;
        }
```

From ‘Missing break in switch’ to Code documentation

- This issue was first detected on 12/21/2016. The importance of commenting ~~your~~ our code.

```
diff --git a/drivers/usb/musb/musb_core.c b/drivers/usb/musb/musb_core.c
index 892088f..d8bae6c 100644
--- a/drivers/usb/musb/musb_core.c
+++ b/drivers/usb/musb/musb_core.c
@@ -1869,6 +1869,7 @@ static void musb_pm_runtime_check_session(struct musb *musb)

                        return;
                }
+               /* fall through */
        case MUSB_QUIRK_A_DISCONNECT_19:
                if (musb->quirk_retries--) {
                        musb_dbg(musb,
```

‘Dereference before null check’

```
diff --git a/drivers/net/ieee802154/ca8210.c b/drivers/net/ieee802154/ca8210.c
index f6df75e..7a21854 100644
--- a/drivers/net/ieee802154/ca8210.c
+++ b/drivers/net/ieee802154/ca8210.c
@@ -912,7 +912,7 @@ static int ca8210_spi_transfer(
 )
 {
        int i, status = 0;
-       struct ca8210_priv *priv = spi_get_drvdata(spi);
+       struct ca8210_priv *priv;
        struct cas_control *cas_ctl;

        if (!spi) {
@@ -923,6 +923,7 @@ static int ca8210_spi_transfer(
                return -ENODEV;
        }

+       priv = spi_get_drvdata(spi);
        reinit_completion(&priv->spi_transfer_complete);

        dev_dbg(&spi->dev, "ca8210_spi_transfer called\n");
```

Fun with NULL pointers (2009)

https://lwn.net/Articles/342330/

From ‘Use after free’ to Code refactoring

- The bug below was first detected on 09/20/2013. Fixed after 4 years.

```
diff --git a/drivers/usb/gadget/udc/amd5536udc.c b/drivers/usb/gadget/udc/amd5536udc.c
index ea03ca7..821d088 100644
--- a/drivers/usb/gadget/udc/amd5536udc.c
+++ b/drivers/usb/gadget/udc/amd5536udc.c
@@ -611,21 +611,20 @@ udc_alloc_request(struct usb_ep *usbep, gfp_t gfp)
 static int udc_free_dma_chain(struct udc *dev, struct udc_request *req)
 {
        int ret_val = 0;
-       struct udc_data_dma     *td;
-       struct udc_data_dma     *td_last = NULL;
+       struct udc_data_dma *td = req->td_data;
        unsigned int i;

+       dma_addr_t addr_next = 0x00;
+       dma_addr_t addr = (dma_addr_t)td->next;
+
        DBG(dev, "free chain req = %p\n", req);

        /* do not free first desc., will be done by free for request */
-       td_last = req->td_data;
-       td = phys_to_virt(td_last->next);
-
        for (i = 1; i < req->chain_len; i++) {
-               pci_pool_free(dev->data_requests, td,
-                               (dma_addr_t)td_last->next);
-               td_last = td;
-               td = phys_to_virt(td_last->next);
+               td = phys_to_virt(addr);
+               addr_next = (dma_addr_t)td->next;
+               pci_pool_free(dev->data_requests, td, addr);
+               addr = addr_next;
        }

        return ret_val;
```

Duplicated code for different branches

- After talking with the maintainer it turned out to be a copy/paste error.

```
diff --git a/drivers/media/platform/qcom/venus/helpers.c b/drivers/media/platform/qcom/venus/helpers.c
index b52410d..68933d2 100644
--- a/drivers/media/platform/qcom/venus/helpers.c
+++ b/drivers/media/platform/qcom/venus/helpers.c
@@ -292,7 +292,7 @@ static void return_buf_error(struct venus_inst *inst,
        if (vbuf->vb2_buf.type == V4L2_BUF_TYPE_VIDEO_OUTPUT_MPLANE)
                v4l2_m2m_src_buf_remove_by_buf(m2m_ctx, vbuf);
        else
-               v4l2_m2m_src_buf_remove_by_buf(m2m_ctx, vbuf);
+               v4l2_m2m_dst_buf_remove_by_buf(m2m_ctx, vbuf);

        v4l2_m2m_buf_done(vbuf, VB2_BUF_STATE_ERROR);
 }
```

Workflow

Workflow and tips

- Review the code around the issue.
- Review it again.
- In case of doubt ask questions to the maintainers.

  (be specific/do your homework).
- Sometimes it is good to ask questions while proposing a patch at the same time.
- Take note of similar cases for future bug fixing.
- Read software security assessment books (sometimes it helps to cope with frustration too).

Problems and the future of this project

- False Positives.
- Use Coccinelle to identify False Positives.
- Continue fixing as much bugs as possible during the next ~8 months.

Contributions

200+ patches upstream

Categories (7)

- NULL pointer dereferences.
- API usage errors.
- Code maintainability issues.
- Control flow issues.
- Uninitialized variables.
- Incorrect expression.
- Integer handling issues.

- Constification (Not a Coverity category)
- Miscellaneous (Not a Coverity category)

Types (21)

- Dereference after null check.
- Dereference before null check.
- Dereference null return value.
- Explicit null dereference.
- Missing null check on return value.
- Arguments in wrong order.
- Ignored error return code.
- Unused value.
- Unused code.
- Unnecessary static on local variable.
- ‘Constant’ variable guards dead code.
- Missing break in switch.
- Uninitialized scalar variable.
- Array compared against 0.
- Identical code for different branches.
- Self assignment.
- Macro compares unsigned to 0.
- Code refactoring.
- Print error message on failure.
- Unnecessary cast on kmalloc.
- Use sizeof(*var) in kmalloc.

Components impacted (26)

- alsa-devel
- ath10k
- dri-devel
- intel-gfx
- linux-arm-kernel
- linux-arm-msm
- linux-clk
- linux-crypto
- linux-fbdev
- linux-fpga
- linux-input
- linux-media
- linux-mediatek
- linux-mmc
- linux-omap
- linux-parisc
- linux-pm
- linux-rdma
- linux-renesas-soc
- linux-rockchip
- linux-scsi
- linux-wireless
- linux-wpan
- platform-driver-x86
- target-devel
- tpmdd-devel

Contributions

Contributor to alsa-devel, ath10k, dri-devel, intel-gfx, linux-amlogic, linux-arm-kernel, linux-arm-msm, linux-block, linux-clk, linux-crypto, linux-dmaengine, linux-fbdev, linux-fpga, linux-input, linux-media, linux-mediatek, linux-mmc, linux-omap, linux-parisc, linux-pci, linux-pm, linux-rdma, linux-renesas-soc, linux-rockchip, linux-samsung-soc, linux-scsi, linux-wireless, linux-wpan, LKML, platform-driver-x86, spi-devel-general, target-devel, tpmdd-devel, xen-devel.

Thank you!

#FuerzaMéxico