Static code checking In the Linux kernel

Presented by

Arnd Bergmann

Date

April 6, 2016

Event

Embedded Linux Conference

Static code checking in the Linux kernel

1. Overview
2. Randconfig issues
3. Checking tools
4. Automated checking

Linaro

Motivation

- Testing arm-soc pull requests
- Testing code refactoring

Linaro

Approaches to build testing

1. Record all known warnings, email about new ones
2. Eliminate all known warnings

Linaro

Timeline for ARM fixes

defconfig failures

clang failures

allmodconfig failures

defconfig warnings

randconfig errors

allmodconfig warnings

randconfig warnings

2011 2012 2013 2014 2015 2016

Linaro

Randconfig issues

Kconfig dependencies

- netfilter
- I2C
- ALSA Codecs
- module, modules, modules

```
--- a/net/openvswitch/Kconfig
+++ b/net/openvswitch/Kconfig
@@ -7,7 +7,9 @@ config OPENVSWITCH
        depends on INET
        depends on !NF_CONNTRACK || \
                   (NF_CONNTRACK && ((!NF_DEFRAG_IPV6 || NF_DEFRAG_IPV6) && \
-                                    (!NF_NAT || NF_NAT)))
+                                     (!NF_NAT || NF_NAT) && \
+                                     (!NF_NAT_IPV4 || NF_NAT_IPV4) && \
+                                     (!NF_NAT_IPV6 || NF_NAT_IPV6)))
        select LIBCRC32C
        select MPLS
        select NET_MPLS_GSO
```

Linaro

Uninitialized variables

- The Power of Undefined Values: http://rusty.ozlabs.org/?p=232

- Problems:
  gcc -Os -fprofile-arcs
  CONFIG_UBSAN_SANITIZE_ALL
  CONFIG_PROFILE_ALL_BRANCHES

Linaro

```
/*
 * "Define 'is'", Bill Clinton
 * "Define 'if'", Steven Rostedt
 */
#define if(cond, ...) __trace_if( (cond , ## __VA_ARGS__) )
#define __trace_if(cond) \
        if (__builtin_constant_p(!!(cond)) ? !!(cond) :                 \
        ({                                                              \
                int ______r;                                            \
                static struct ftrace_branch_data                        \
                        __attribute__((__aligned__(4)))                 \
                        __attribute__((section("_ftrace_branch")))      \
                        ______f = {                                     \
                                .func = __func__,                       \
                                .file = __FILE__,                       \
                                .line = __LINE__,                       \
                        };                                              \
                ______r = !!(cond);                                     \
                ______f.miss_hit[______r]++;                            \
                ______r;                                                \
        }))
```

```
        if (cm_node->ipv4)
                arpindex = i40iw_addr_resolve_neigh(iwdev,
                                cm_info->loc_addr[0], cm_info->rem_addr[0],
                                oldarpindex);
#if IS_ENABLED(CONFIG_IPV6)
        else
                arpindex = i40iw_addr_resolve_neigh_ipv6(iwdev,
                                cm_info->loc_addr, cm_info->rem_addr,
                                oldarpindex);
#endif

        ether_addr_copy(cm_node->rem_mac, iwdev->arp_table[arpindex].mac_addr);
```

```
if (cm_node->ipv4)
       arpindex = i40iw_addr_resolve_neigh(iwdev,
                       cm_info->loc_addr[0], cm_info->rem_addr[0],
                       oldarpindex);
else if (CONFIG_IPV6)
      arpindex = i40iw_addr_resolve_neigh_ipv6(iwdev,
                       cm_info->loc_addr, cm_info->rem_addr,
                       oldarpindex);
else
        arpindex = -EINVAL;
ether_addr_copy(cm_node->rem_mac, iwdev->arp_table[arpindex].mac_addr);
```

Linaro

```
@@ -425,8 +425,8 @@ static int br_mdb_add_group(struct net_bridge *br,
struct net_bridge_port *port,
        mp = br_mdb_ip_get(mdb, group);
        if (!mp) {
                mp = br_multicast_new_group(br, port, group);
-               err = PTR_ERR(mp);
-               if (IS_ERR(mp))
+               err = PTR_ERR_OR_ZERO(mp);
+               if (err)
                        return err;
        }
```

Checking tools

scripts/checkpatch.pl

- Written by Andy Whitcroft, Joe Perches
- String matching
- Checks for basic coding style issues
- Good for checking new submissions, less so for existing code

Linaro

```
WARNING: line over 80 characters
#56: FILE: kernel/sched/sched.h:56:
+#if 0 /* BITS_PER_LONG > 32 -- currently broken: it increases power usage under
WARNING: Unnecessary space before function pointer arguments
#1187: FILE: kernel/sched/sched.h:1187:
+       void (*enqueue_task) (struct rq *rq, struct task_struct *p, int flags);
WARNING: please, no spaces at the start of a line
#1252: FILE: kernel/sched/sched.h:1252:
+   for (class = sched_class_highest; class; class = class->next)$
WARNING: Prefer 'unsigned int' to bare use of 'unsigned'
#1346: FILE: kernel/sched/sched.h:1346:
+static inline void add_nr_running(struct rq *rq, unsigned count)
WARNING: please, no spaces at the start of a line
#1816: FILE: kernel/sched/sched.h:1816:
+       if (data)$
WARNING: suspect code indent for conditional statements (7, 15)
#1816: FILE: kernel/sched/sched.h:1816:
+       if (data)
+               data->func(data, time, util, max);
```

sparse



- Domain specific checks
- Written by Linus Torvalds for use with Linux
- Later maintained by Josh Triplett,Chris Li

```
make C=1 CHECK="/usr/bin/sparse"
```

http://git.kernel.org/pub/scm/devel/sparse/chrisl/sparse.git

Linaro

```
kernel/sys.c:948:32: warning: incorrect type in argument 1 (different address
spaces)
kernel/sys.c:948:32:    expected struct task_struct *p1
kernel/sys.c:948:32:    got struct task_struct [noderef] <asn:4>*real_parent
kernel/fork.c:1231:41: warning: implicit cast to nocast type
kernel/fork.c:1324:13: warning: dereference of noderef expression
kernel/irq/irqdesc.c:567:17: warning: context imbalance in
'__irq_get_desc_lock' - wrong count at exit
```

Linaro

Extra gcc warnings

- `make W=1`
  Added by Borislav Petkov
  Generally useful warnings
- `make W=12`
  Possibly useful warnings
- `make W=123`
  overload

Linaro

```
arch/arm/mach-tegra/cpuidle-tegra20.c:216:6: warning: no previous
                 prototype for 'tegra20_cpuidle_pcie_irqs_in_use'
                 [-Wmissing-prototypes]
mm/vmscan.c: In function 'try_to_free_mem_cgroup_pages':
mm/vmscan.c:2907:6: warning: variable 'nid' set but not used
                           [-Wunused-but-set-variable]
security/keys/trusted.c: In function 'tpm_unseal':
security/keys/trusted.c:589:11: warning: variable 'keyhndl'
                     set but not used [-Wunused-but-set-variable]
drivers/pwm/pwm-bcm-kona.c: In function 'kona_pwmc_config':
drivers/pwm/pwm-bcm-kona.c:141:35: warning: comparison of unsigned
                          expression < 0 is always false [-Wtype-limits]
```

Linaro

Extra gcc warnings

```
$ make W=1 2>&1 | cut -f 2 -d[ | sort | uniq -c | cut -f 1 -d] | sort -nr

   631 -Woverride-init
   280 -Wmissing-prototypes
   133 -Wunused-but-set-variable
    98 -Wmissing-include-dirs
    33 -Wtype-limits
    13 -Wsuggest-attribute=format
    13 -Wempty-body
     7 -Wold-style-declaration
     2 -Wignored-qualifiers
```

Linaro

Extra gcc warnings

```
$ make W=12 2>&1 | cut -f 2 -d[ | sort | uniq -c | cut -f 1 -d] | sort -nr

 94235 -Wnested-externs
 93350 -Wcast-align
 16694 -Wsign-compare
  7594 -Wshadow
  1465 -Wmissing-field-initializers
   631 -Woverride-init
   482 -Waggregate-return
   280 -Wmissing-prototypes
   133 -Wunused-but-set-variable
    98 -Wmissing-include-dirs
    33 -Wtype-limits
    13 -Wsuggest-attribute=format
    13 -Wempty-body
       ...
```

Linaro

Extra gcc warnings

```
$ make W=123 2>&1 | cut -f 2 -d[ | sort | uniq -c | cut -f 1 -d] | sort -nr

782719 -Wsign-conversion
425231 -Wpadded
176825 -Wcast-qual
176553 -Wconversion
 94235 -Wnested-externs
 93350 -Wcast-align
 56153 -Wpointer-arith
 45309 -Wredundant-decls
 36921 -Wbad-function-cast
 27978 -Wattributes
 16694 -Wsign-compare
 13440 -Wpacked
  7594 -Wshadow
  6820 -Wswitch-default
       ...
```

Linaro

gcc-6 warnings

- Overall helpful additions
- 32 patches so far, most applied

https://gnu.wildebeest.org/blog/mjw/2016/02/15/looking-forward-to-gcc6-many-new-warnings/

Linaro

```
@@ -2860,7 +2860,7 @@ lpfc_online(struct lpfc_hba *phba)
         }

         vports = lpfc_create_vport_work_array(phba);
-        if (vports != NULL)
+        if (vports != NULL) {
                 for (i = 0; i <= phba->max_vports && vports[i] != NULL; i++) {
                         struct Scsi_Host *shost;
                         shost = lpfc_shost_from_vport(vports[i]);
@@ -2877,7 +2877,8 @@ lpfc_online(struct lpfc_hba *phba)
                         }
                         spin_unlock_irq(shost->host_lock);
                 }
-                lpfc_destroy_vport_work_array(phba, vports);
+        }
+        lpfc_destroy_vport_work_array(phba, vports);

         lpfc_unblock_mgmt_io(phba);
         return 0;
```

Linaro

llvmlinux

- Building the kernel with clang
- Patches needed to just compile
- Tons of new warnings
  some good, some less so
- Lots of work done by Behan Webster

http://git.kernel.org/cgit/linux/kernel/git/arnd/playground.git/log/?h=llvmlinux-4.5

Linaro

```
static int __init test_static_key_init(void)
{
#define test_key_func(key, branch) \
      {bool func(void) { return branch(key); } func; })

      struct test_key static_key_tests[] = {
            {
                  .init_state  = true,
                  .key         = &old_true_key,
                  .test_key    = test_key_func(&old_true_key, static_key_true),
             },
             {
                  .init_state  = false,
                  .key         = &old_false_key,
                  .test_key    = test_key_func(&old_false_key, static_key_false),
             },
             ...
      };
      ...
}
```

Linaro

llvmlinux

- Picked up the 3.19 based tree in 2016
- Ported to v4.5
- builds ARM randconfig successfully now
- TODO: address warnings

```
drivers/staging/wilc1000/wilc_spi.c:123:34: warning: tentative definition
of variable with internal linkage has incomplete non-array type 'const
struct wilc1000_ops' [-Wtentative-definition-incomplete-type]

--- a/drivers/staging/wilc1000/wilc_spi.c
+++ b/drivers/staging/wilc1000/wilc_spi.c
@@ -120,8 +120,6 @@ static u8 crc7(u8 crc, const u8 *buffer, u32 len)

 #define USE_SPI_DMA     0

-static const struct wilc1000_ops wilc1000_spi_ops;
-
 static int wilc_bus_probe(struct spi_device *spi)
 {
        int ret, gpio;
```

Linaro

```
3502 -Wgnu-variable-sized-type-not-at-end
 665 -Winitializer-overrides
 628 -Wduplicate-decl-specifier
 542 -Wbitfield-constant-conversion
  42 -Wunused-const-variable
  24 -Wenum-conversion
  22 -Wtautological-compare
  14 -Wformat-invalid-specifier
  10 -Wtautological-constant-out-of-range-compare
  10 -Wsometimes-uninitialized
  10 -Wformat
   7 -Wshift-negative-value
   4 -Wpointer-bool-conversion
   4 -Wdeprecated-declarations
   3 -Wunused-value
   3 -Wunneeded-internal-declaration
   3 -Wuninitialized
   3 -Wshift-count-overflow
   3 -Wframe-larger-than=
   2 -Wtautological-pointer-compare
     ...
```

Linaro

```
make -skj20 CC=/usr/bin/clang 2>&1 | cut -f 2 -d[ | sort | uniq
-c | cut -f 1 -d] | sort -nr

    542 -Wbitfield-constant-conversion
     24 -Wenum-conversion
      7 -Wshift-negative-value
      4 -Wpointer-bool-conversion
      4 -Wdeprecated-declarations
      3 -Wunneeded-internal-declaration
      3 -Wshift-count-overflow
      3 -Wframe-larger-than=
      2 -Wself-assign
      2 -Wconstant-conversion
      1 -Wsection
      1 -Warray-bounds
```

Linaro

```
make -skj20 CC=/usr/bin/clang 2>&1 | cut -f 2 -d[ | sort | uniq
-c | cut -f 1 -d] | sort -nr

    542 -Wbitfield-constant-conversion
     24 -Wenum-conversion
      7 -Wshift-negative-value
      4 -Wpointer-bool-conversion
      4 -Wdeprecated-declarations
      3 -Wunneeded-internal-declaration
      3 -Wshift-count-overflow
      3 -Wframe-larger-than=
      2 -Wself-assign
      2 -Wconstant-conversion
      1 -Wsection
      1 -Warray-bounds
```

all in one file!

Linaro

clang static analyser

- Additional checks on top of llvmlinux
- Local http interface
- Allows defining domain specific checks

http://linuxplumbersconf.org/2015/ocw//system/presentations/3237/original/2015-LPC-Clang-static-analyzer-kernel.pdf

Linaro

```
1932  void spi_unregister_master(struct spi_master *master)
1933  {
1934          int dummy;
1935
1936          if (master->queued) {
1937                  if (spi_destroy_queue(master))
1938                          dev_err(&master->dev, "queue remove failed\n");
1939          }
1940
1941          mutex_lock(&board_lock);
1942          list_del(&master->list);
1943          mutex_unlock(&board_lock);
1944
1945          dummy = device_for_each_child(&master->dev, NULL, __unregister);
                  Value stored to 'dummy' is never read
1946          device_unregister(&master->dev);
1947  }
1948  EXPORT_SYMBOL_GPL
      extern typeof(spi_unregister_master) spi_unregister_master;
       void *  __crc_spi_unregister_master   __attribute__  ((weak)); s
```

Linaro

Synopsys Coverity

- Commercial code checker
- Focus on security bugs
- x86 only (use COMPILE_TEST!)
- Requires manual categorization of bugs
- Lots of work done by Dave Jones in the past

https://scan.coverity.com/projects/linux

http://codemonkey.org.uk/2014/08/13/year-coverity-linux-kernel-scans

Linaro

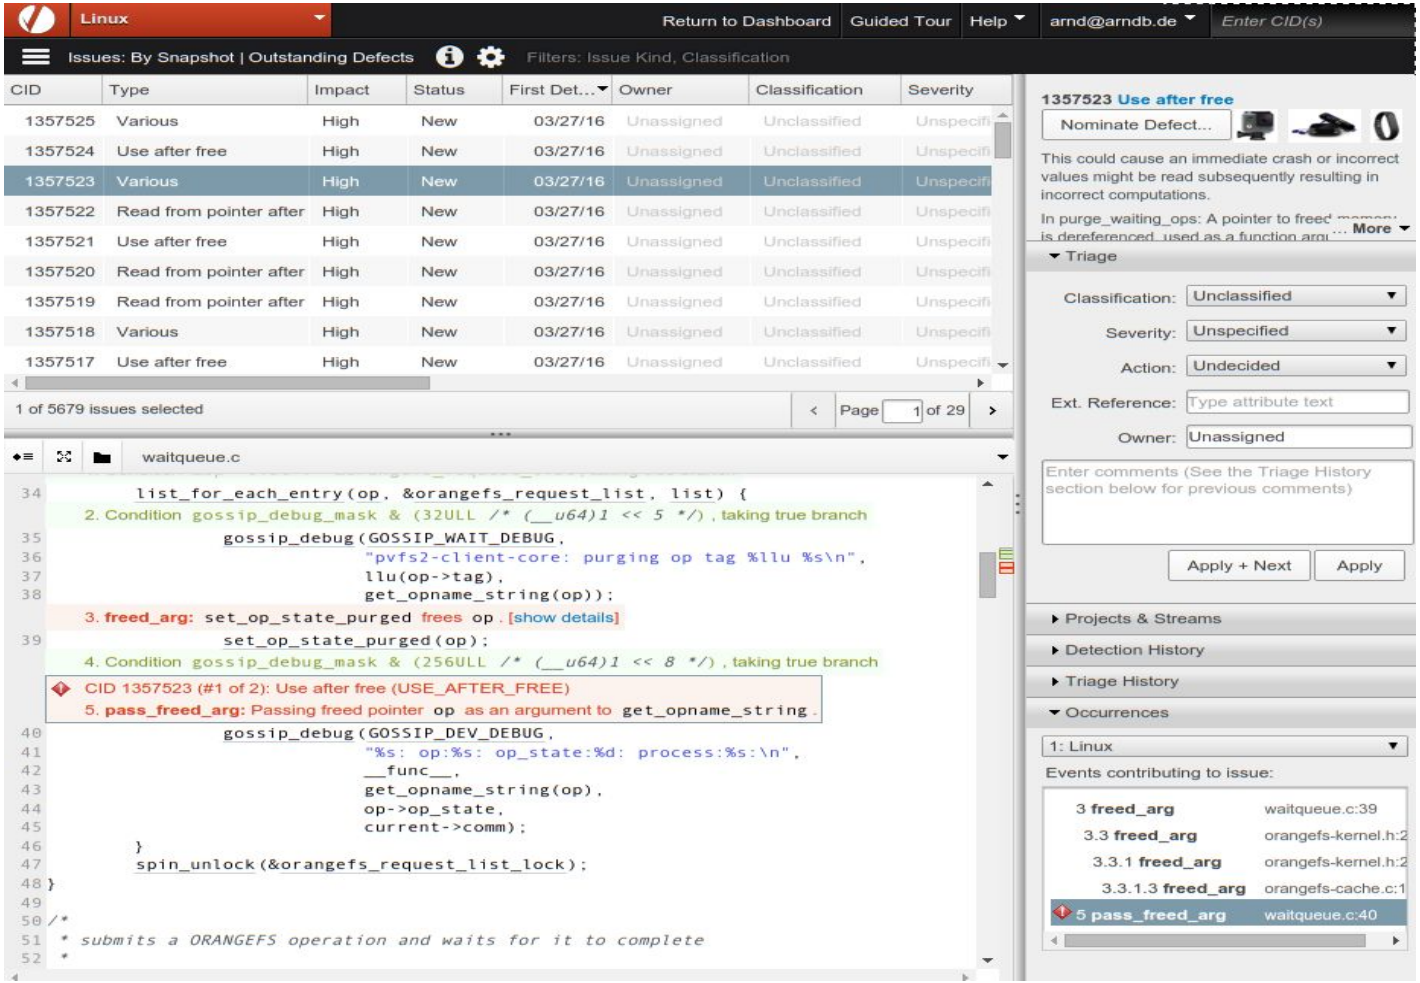
Linux
Return to Dashboard
Guided Tour
Help
arnd@arndb.de
Enter CID(s)
Issues: By Snapshot | Outstanding Defects
Filters: Issue Kind, Classification
CID
Type
Impact
Status
First Det...
Owner
Classification
Severity
1357525 Various High New 03/27/16 Unassigned Unclassified Unspecifi
1357524 Use after free High New 03/27/16 Unassigned Unclassified Unspecifi
1357523 Various High New 03/27/16 Unassigned Unclassified Unspecifi
1357522 Read from pointer after High New 03/27/16 Unassigned Unclassified Unspecifi
1357521 Use after free High New 03/27/16 Unassigned Unclassified Unspecifi
1357520 Read from pointer after High New 03/27/16 Unassigned Unclassified Unspecifi
1357519 Read from pointer after High New 03/27/16 Unassigned Unclassified Unspecifi
1357518 Various High New 03/27/16 Unassigned Unclassified Unspecifi
1357517 Use after free High New 03/27/16 Unassigned Unclassified Unspecifi
1 of 5679 issues selected
Page 1 of 29
waitqueue.c
list_for_each_entry(op, &orangefs_request_list, list) {
2. Condition gossip_debug_mask & (32ULL /* (__u64)1 << 5 */), taking true branch
gossip_debug(GOSSIP_WAIT_DEBUG,
"pvfs2-client-core: purging op tag %llu %s\n",
llu(op->tag),
get_opname_string(op));
3. freed_arg: set_op_state_purged frees op. [show details]
set_op_state_purged(op);
4. Condition gossip_debug_mask & (256ULL /* (__u64)1 << 8 */), taking true branch
CID 1357523 (#1 of 2): Use after free (USE_AFTER_FREE)
5. pass_freed_arg: Passing freed pointer op as an argument to get_opname_string.
gossip_debug(GOSSIP_DEV_DEBUG,
"%s: op:%s: op_state:%d: process:%s:\n",
__func__,
get_opname_string(op),
op->op_state,
current->comm);
}
spin_unlock(&orangefs_request_list_lock);
}
/*
* submits a ORANGEFS operation and waits for it to complete
*
1357523 Use after free
Nominate Defect...
This could cause an immediate crash or incorrect values might be read subsequently resulting in incorrect computations.
In purge_waiting_ops: A pointer to freed
More
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
Events contributing to issue:
3 freed_arg waitqueue.c:39
3.3 freed_arg orangefs-kernel.h:2
3.3.1 freed_arg orangefs-kernel.h:2
3.3.1.3 freed_arg orangefs-cache.c:1
5 pass_freed_arg waitqueue.c:40

Linaro

Coccinelle

- Written by Julia Lawall
- http://coccinelle.lip6.fr/
- Can patch code or just warn about it
- Very sophisticated pattern matching

```
make C=1 CHECK="scripts/coccicheck"
```

Linaro

```
arch/arm/mm/dma-mapping.c:839:36-42: WARNING: Consider using vma_pages helper
on vma
init/do_mounts_initrd.c:129:40-48: Move constant to right.
init/calibrate.c:31:36-38: Move constant to right.
arch/arm/mm/fault.c:533:2-5: WARNING: Use BUG_ON instead of if condition
followed by BUG.
Please make sure the condition has no side effects (see conditional BUG_ON
definition in include/asm-generic/bug.h)
arch/arm/kernel/module.c:379:2-44: code aligned with following code on line
381
arch/arm/mach-artpec/board-artpec6.c:46:2-3: Unneeded semicolon
arch/arm/kernel/topology.c:99:1-15: alloc with no test, possible model on line
107
```

```
--- a/drivers/staging/rdma/hfi1/chip.c
+++ b/drivers/staging/rdma/hfi1/chip.c
@@ -1250,11 +1250,8 @@ CNTR_ELEM(#name, \

 u64 read_csr(const struct hfi1_devdata *dd, u32 offset)
 {
-       u64 val;
-
        if (dd->flags & HFI1_PRESENT) {
-               val = readq((void __iomem *)dd->kregbase + offset);
-               return val;
+               return readq((void __iomem *)dd->kregbase + offset);
        }
        return -1;
 }
```

Linaro

smatch

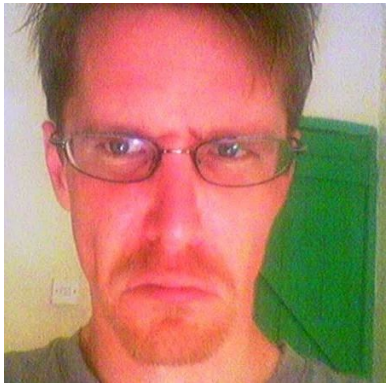

- Written by Dan Carpenter
- 3000 bugs fixed so far (mostly by Dan)

https://blogs.oracle.com/linuxkernel/entry/smatch_static_analysis_tool_overview

http://repo.or.cz/w/smatch.git

```
make C=1 CHECK="smatch -p=kernel"
```

Linaro

```
drivers/gpu/drm/msm/mdp/mdp5/mdp5_plane.c:198 mdp5_plane_reset() error:
    potential null dereference 'mdp5_state'.  (kzalloc returns null)
drivers/gpu/drm/msm/mdp/mdp5/mdp5_plane.c:225 mdp5_plane_duplicate_state()
     error: we previously assumed 'mdp5_state' could be null (see line 222)
drivers/gpu/drm/msm/mdp/mdp5/mdp5.xml.h:554 __offset_PIPE() info: ignoring
     unreachable code.
drivers/gpu/drm/msm/mdp/mdp5/mdp5.xml.h:943 __offset_SW_PIX_EXT() info:
     ignoring unreachable code.
drivers/clk/clk-gpio.c:132 clk_register_gpio() warn: possible memory
    leak of 'clk_gpio'
include/linux/skbuff.h:2404 __netdev_alloc_skb_ip_align() warn:
    should this be a bitwise op?
drivers/mmc/host/dw_mmc.c:1794 dw_mci_tasklet_func() warn: missing break?
     reassigning 'state'
drivers/net/ethernet/renesas/sh_eth.c:2135 sh_eth_get_strings() error:
     memcpy() '*sh_eth_gstrings_stats' too small (32 vs 128)
drivers/net/ethernet/realtek/r8169.c:2320 rtl8169_get_strings() error:
     memcpy() '*rtl8169_gstrings' too small (32 vs 416)
```

Automated checking

0day build bot

- Maintained by Fengguang Wu
- Tests public git trees
- Now tests patch submissions
- Automatic bisection of bugs
- Even sends patches

Linaro

```
Re: [PATCH] mtd: avoid stack overflow in MTD CFI code
From:    kbuild test robot <lkp@intel.com>

Hi Arnd,

[auto build test WARNING on v4.5-rc1]
[also build test WARNING on next-20160125]
[if your patch is applied to the wrong git tree, please drop us a note to help
improving the system]

url:    https://github.com/0day-ci/linux/commits/Arnd-Bergmann/mtd-avoid-stack-
overflow-in-MTD-CFI-code/20160125-234611
config: um-allmodconfig (attached as .config)
reproduce:
        # save the attached .config to linux build tree
        make ARCH=um

All warnings (new ones prefixed by >>):

warning: (MTD_MAP_BANK_WIDTH_32) selects MTD_COMPLEX_MAPPINGS which has unmet
direct dependencies (MTD && HAS_IOMEM)
```

```
commit e014e8468552236f0f0cb64c7c341c1dce506070
Author: Wu Fengguang <fengguang.wu@intel.com>
Date:   Sat Mar 19 00:54:50 2016 +0800

    ovs: internal_set_rx_headroom() can be static

    Signed-off-by: Fengguang Wu <fengguang.wu@intel.com>
    Signed-off-by: David S. Miller <davem@davemloft.net>

--- a/net/openvswitch/vport-internal_dev.c
+++ b/net/openvswitch/vport-internal_dev.c
@@ -138,7 +138,7 @@ internal_get_stats(struct net_device *dev, struct
rtnl_link_stats64 *stats)
        return stats;
 }

-void internal_set_rx_headroom(struct net_device *dev, int new_hr)
+static void internal_set_rx_headroom(struct net_device *dev, int new_hr)
 {
        dev->needed_headroom = new_hr;
 }
```

kernelci.org

- Build and boot testing
- Focused on ARM machines
- Builds defconfigs, allmodconfig
- Distributed infrastructure

https://fosdem.org/2016/schedule/event/kernelci/attachments/slides/888/export/events/attachments/kernelci/slides/888/kernelci_org__FOSDEM_2016.pdf

Linaro

Home | Jobs | Builds | Boots | SoCs | Compare β | Info

Available Jobs

The results shown here cover the last **14 days** of available data starting from **Wed, 30 Mar 2016** (time is UTC based).

25 ▾ **jobs per page**

Filter the results

Tree	Branch	Latest Build Status	Latest Boot Status	Date	Status
broonie-regmap	local/for-next	136 134 2 0	193 162 16 15	2016-03-30	
broonie-regulator	local/for-next	137 135 2 0	36 33 3 0	2016-03-30	
lsk	local/linux-linaro-lsk-v4.1-test	141 141 0 0	496 464 16 16	2016-03-30	
next	local/master	137 133 4 0	455 386 54 15	2016-03-30	
broonie-sound	local/for-next	137 135 2 0	204 186 4 14	2016-03-29	
evalenti	local/for-kernelci	137 135 2 0	476 439 19 18	2016-03-29	
stable-sasha	local/linux-3.18.y-queue	149 149 0 0	360 336 10 14	2016-03-29	
ulfh	local/next	137 135 2 0	456 419 19 18	2016-03-29	
broonie-spi	local/for-next	137 135 2 0	185 166 5 14	2016-03-29	
mainline	local/master	137 135 2 0	427 388 20 19	2016-03-28	
renesas	local/devel	137 135 2 0	449 404 26 19	2016-03-28	
pmwg	local/integ	136 133 3 0	444 418 21 5	2016-03-22	

Linaro

Questions?