The ISoLA 2012 symposium Heraklion, Crete 15 Oct 2012 Linux Driver Verification Track



Linux Device-Drivers Verification Challenges

<u>Alexey Khoroshilov</u> Vadim Mutilin Eugene Novikov



Institute for System Programming of the Russian Academy of Sciences



Definitions to discuss



Static Analysis: Trade-Off Triangle



Static Analysis vs Model Checking





LinuxTesting •Org

Linux Devic

Large and constant

- important —

6,5 MI

(*) drivers&sc

- not too big
- not too compl

software code

Source code available

Sounds like a dream



But...



But...

(1) No entry point



Reachability problem

LinuxTesting

ord



error location



module_init(DAC960_init_module);
module_exit(DAC960_cleanup_module);

LinuxTesting

ord

procedures

Linux Kernel



(1) No entry point

- No predefined entry point
- Event handlers via function pointers
- Order limitation
 - open() after probe(), but before remove()
- Implicit limitations
 - read() only if open() succeed
- and the limitations are specific for each class of drivers



But...

(1) No entry point

(2) Kernel core model

Linux Kernel





(2) Kernel core model

- Boundary of driver
- Kernel core API properties



But...

(1) No entry point

(3) No stable API

(2) Kernel core model



This is being written to try to explain why Linux does not have a binary kernel interface, nor does it have a stable kernel interface. Please realize that this article describes the **in kernel** interfaces, not the kernel to userspace interfaces. The kernel to userspace interface is the one that application programs use, the syscall interface. That interface is **very** stable over time, and will not break.

Executive Summary

LinuxTesting

orq

You think you want a stable kernel interface, but you really do not, and you don't even know it. What you want is a stable running driver, and you get that only if your driver is in the main kernel tree. You also get lots of other good benefits if your driver is in the main kernel tree, all of which has made Linux into such a strong, stable, and mature operating system which is the reason you are using it in the first place.

Greg Kroah-Hartman



(3) No stable API

- Environment interface and invariants
- Kernel core model

But...

ord

LinuxTesting

(1) No entry point

(3) No stable API

(2) Kernel core model

(4) Low level code



container_of

#define container_of(ptr, type, member) ({ const typeof(((type *)0) \rightarrow member) * __mptr = (ptr); (type *)((char *) __mptr - offsetof(type,member));})



(4) Low level code

- pointer arithmetics
- casting
- container_of

LinuxTesting •Org

But...

(1) No entry point

(3) No stable API

(2) Kernel core model

(4) Low level code

(5) Hardware specific

Linux Kernel





(5) Hardware specific

- Hardware specific bugs
- Hardware specific invariants

LinuxTesting •Org

But...

(1) No entry point

(3) No stable API

(2) Kernel core model

(4) Low level code

(5) Hardware specific

(6) Concurrency

LinuxTesting

- Device drivers are significantly asynchronous
- Many bugs appears in concurrent settings only

Commit Analysis

- All patches in stable trees (2.6.35 3.0) for 1 year:
 - 26 Oct 2010 26 Oct 2011
- 1503 patches in device drivers
- Main goal: detect and classify typical bugs

Commit Analysis (2)

LinuxTesting

orq

	Class	#	%		Class	#	%
1	sync:race	60	17.2%	12	specific:net	10	2.9%
2	specific:resource	32	9.2%	13	specific:usb	9	2.6%
3	generic:null_ptr_deref	31	8.9%	14	generic:int_overflow	8	2.3%
4	specific:check_params	25	7.2%	15	generic:buffer_overflow	8	2.3%
5	generic:resource	24	6.9%	16	specific:check_ret_val	7	2.0%
6	specific:context	19	5.4%	17	generic:uninit	6	1.7%
7	specific:uninit	17	4.9%	10	• • • •	4	1 10/
8	generic:syntax	14	4.0%	18	specific:dma	4	1.1%
9	specific:lock	12	3.4%	19	specific:device	4	1.1%
10	sync:deadlock	11	3.2%	20	specific:misc	27	7.7%
11	spacific:style	10	7 0%	21	generic:misc	11	3.2%
TT	specific.style	10	2.570				



(1) No entry point

(3) No stable API

(2) Kernel core model

(4) Low level code

(5) Hardware specific

Institute for System Programming of Russian Academy of Sciences

VERIFICATION CENTER LINUX



founded in 2005

LinuxTesting

ord

- OLVER Program
- Linux Standard Base Infrastructure Program
- Linux Driver Verification Program



(1) No entry point

(3) No stable API

(5) Hardware specific

(2) Kernel core model

(4) Low level code

Pointer Analysis with Uninterpreted Functions

M. Mandrykin

(1) No entry point

(3) No stable API

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(5) Hardware specific

(2) Kernel core model

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(4) Low level code

Pointer Analysis with Uninterpreted Functions

M. Mandrykin

(1) No entry point

Environment model generator

(3) No stable API

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(5) Hardware specific

(2) Kernel core model

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(4) Low level code

Pointer Analysis with Uninterpreted Functions

M. Mandrykin

```
LinuxTesting
•Org
```

```
Pseudo-main generation
int main(int argc, char* argv[])
{
  init_module()
  for(;;) {
    switch(*) {
     case 0: driver_probe(*,*,*);break;
     case 1: driver_open(*,*);break;
     . . .
 exit_module();
```



(1) No entry point

Environment model generator

(3) No stable API

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(5) Hardware specific

(2) Kernel core model

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(4) Low level code

Pointer Analysis with Uninterpreted Functions

M. Mandrykin

LinuxTesting •Org

Future Data Race Detection Tools

- Runtime analysis
 - Kernel Strider (Google Research Award)
 - KEDR(ISPRAS) + ThreadSaniziter(Google)
 - Race Hound
 - HW breakpoints

Beta to be released by the end of 2012 for x86 only

- Static analysis
 - Research in progress

Conclusions

(1) No entry point

Environment model generator

(3) No stable API

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(5) Hardware specific _{No idea}

(2) Kernel core model

Using Aspect-Oriented Programming for Preparing C Programs for Static Verification

E. Novikov

(4) Low level code

Pointer Analysis with Uninterpreted Functions

M. Mandrykin

(6) Concurrency
To be researched



Still positive



Crete, 2005

LinuxTesting Org

Thank you!

Alexey Khoroshilov khoroshilov@linuxtesting.org http://linuxtesting.org/project/ldv



Institute for System Programming of the Russian Academy of Sciences