

Effectiveness of Driver Isolation and Testing in User Space



TECHNISCHE
UNIVERSITÄT
DARMSTADT

11th EuroSys Doctoral Workshop (EuroDW 2017), April 23, 2017, Belgrade, Serbia

```
... linked in: usb_storage uhci_hcd sd_mod scsi_mod tun ide_cd cdrom ceph100 via_r...
C0160894 LR: C01609B8 CTR: C0160AF4
ca2b7960 TRAP: 0700 Not tainted (2.6.18.1-erik)
00021032 <ME,IR,DR> CR: 82202488 XER: 00000000
l = 4e0ec7b0(9823) 'dangsguardian' THREAD: ca2b6000
00: 00000000 CA2B7A60 DE0EC7B0 00000F22 CE9FES24 0000000B EF921DD2 00000000
08: CFD47A80 00000F23 00001000 00000001 2793349D 18077880 77C6C0F8 00000001
16: 00000000 CA2B7EE0 000005B4 000005B4 000005B4 77C6C0F0 00000000 00000003
24: 00000A4C EF9E98A0 00000000 80600000 CE9FES24 CFD47A80 CE9FES24 CFD47800
P (C0160894) eth_alloc_tx_desc_index+0x48<0x50
(C01609B8) eth_tx_submit_descs_for...
all Trace:
CA2B7A60] [CA2B7A50] 0xca2b7a50 (unre
CA2B7A90] [C016089C] m643xx_eth_star
CA2B7AC0] [C020300C] dev_hard_start_x
CA2B7AE0] [C02032A8] dev_queue_xmit+0
CA2B7B00] [C0244A4C] ip_output+0x198/
CA2B7B30] [C0244F5C] ip_queue_xmit+0x
CA2B7C20] [C0257608] tcp_transmit_skb
CA2B7C60] [C0259228] tcp_push_one+0x9
CA2B7C80] [C024C748] tcp_sendmsg+0x3d
CA2B7D00] [C0268330] inet_sendmsg+0xd
CA2B7D20] [C01F6734] sock_sendmsg+0xb
CA2B7E20] [C01F7D88] sys_sendto+0xd4/
CA2B7F00] [C01F8694] sys_socketcall+0
CA2B7F40] [C001260C] ret_from_syscal
--- Exception: c01 at 0xf8b0e5c
LR = 0x100120a8
Instruction dump:
39200000 0f090000 80680020 81680024 3
742b5a78 200b0000 74605914 91280020 <0f0b0000> 4e800020 9421ff40 7c0002a6
<0>Kernel panic - not syncing: Fatal exception in interrupt
<0>Rebooting in 100 seconds...
```

Oliver Schwahn
os@cs.tu-darmstadt.de

PhD Advisor: Prof. Neeraj Suri
DEEDS Group
Department of Computer Science
TU Darmstadt, Germany

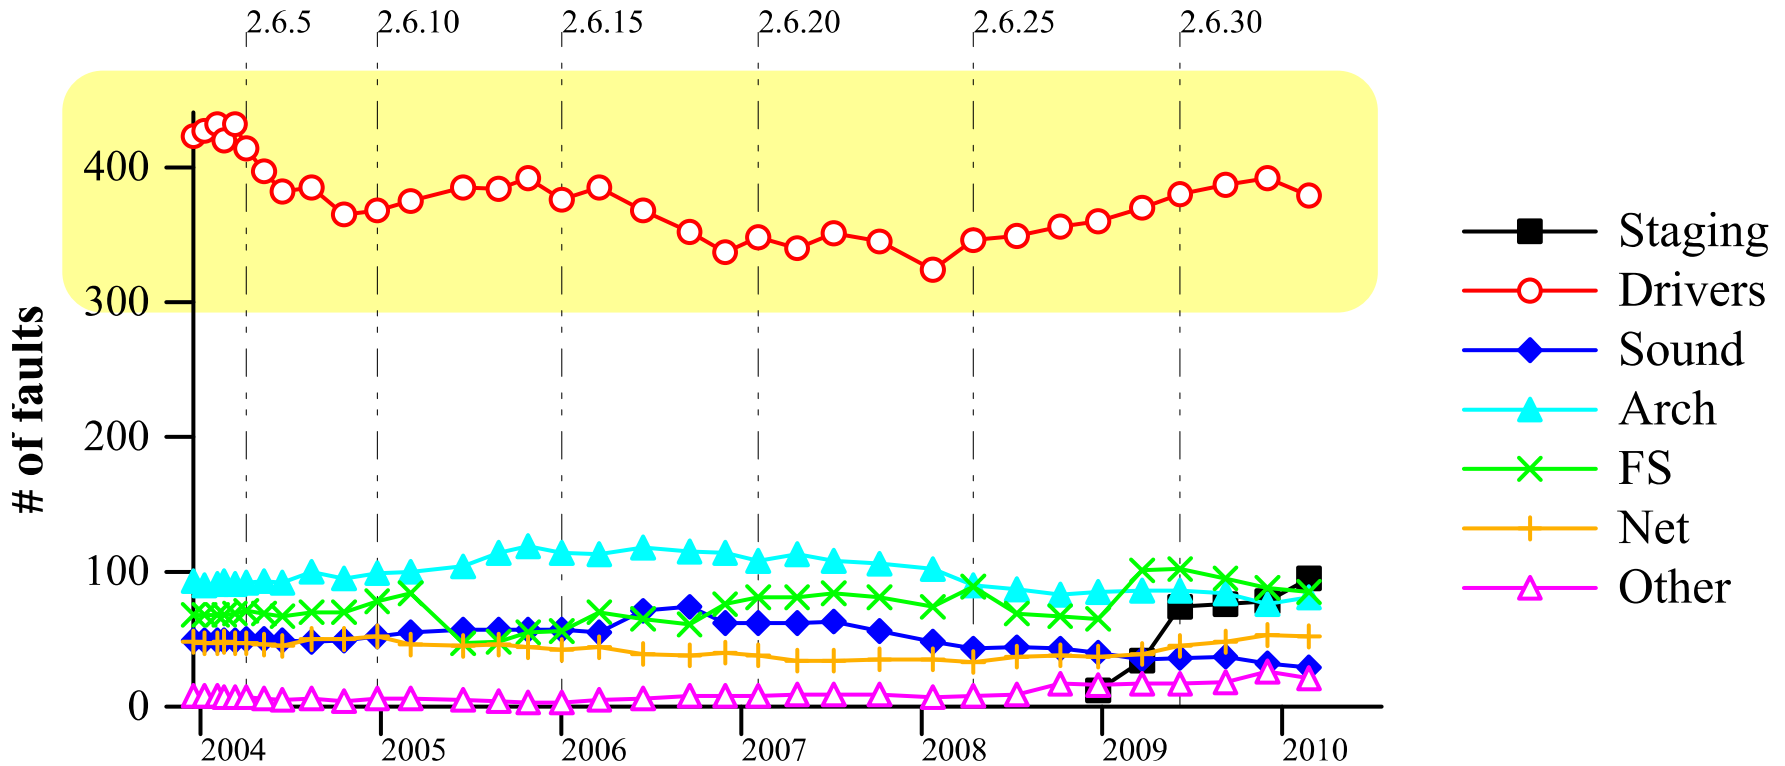
www.deeds.informatik.tu-darmstadt.de



```
uhci_hcd ata_piix libata ehci_hcd scsi_mod usbcore bnx2 nls_base [last unloaded: scsi_wait_scan]
[144516.020026]
[144516.020026] Pid: 0, comm: kworker/0:0 Not tainted 2.6.38-1-amd64 #1 IBM IBM eServer BladeCenter HS21 -[8B
[144516.020026] RIP: 0010:[<ffffffff812971de>] [<ffffffff812971de>] ip_options_compile+0x1c0/0x41a
[144516.020026] RSP: 0018:ffff8800cfc83bc0  EFLAGS: 00010286
[144516.020026] RAX: 0000000000000008 RBX: ffff8800a5aed100 RCX: ffff8801282e2077
[144516.020026] RDX: 000000000000000b RSI: 0000000000000134 RDI: ffffffff8180b500
[144516.020026] RBP: ffff8801282e2075 R08: ffffffff8041bdbb R09: ffff8800cfc83d18
[144516.020026] R10: ffffffff8104c5a3 R11: dead000000200200 R12: ffff8800a5aed128
[144516.020026] R13: 0000000000000027 R14: ffff8801282e2060 R15: 0000000000000027
[144516.020026] FS: 0000000000000000(0000) GS:ffff8800cfc80000(0000) knlGS:0000000000000000
[144516.020026] CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b
[144516.020026] CR2: 0000000000000134 CR3: 00000001284ac000 CR4: 000000000000006e0
[144516.020026] DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000
[144516.020026] DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7: 00000000000000400
[144516.020026] Process kworker/0:0 (pid: 0, threadinfo ffff88012b6d2000, task ffff88012b65ca40)
[144516.020026] Stack:
[144516.020026] ffff880180000000 ffff8800cfc83bd0 0000000000000134 ffffffff8180b500
[144516.020026] ffff880129a4a900 ffffffff80420dab ffffffff00000001 ffff8800a5aed128
[144516.020026] ffff8800a5aed100 ffff880128084000 ffff8801282e2060 ffff880128084000
[144516.020026] Call Trace:
[144516.020026] <IRQ>
[144516.020026] [<fffffff80420dab>] ? br_nf_pre_routing_finish+0x236/0x258 [bridge]
[144516.020026] [<fffffff8041ff36>] ? br_parse_ip_options+0x131/0x19c [bridge]
[144516.020026] [<fffffff80420a6b>] ? br_nf_pre_routing+0x410/0x51a [bridge]
[144516.020026] [<fffffff8128f304>] ? nf_iterate+0x41/0x7e
[144516.020026] [<fffffff8041bdbb>] ? br_handle_frame_finish+0x0/0x1c4 [bridge]
[144516.020026] [<fffffff8128f3a3>] ? nf_hook_slow+0x62/0xde
[144516.020026] [<fffffff8041bdbb>] ? br_handle_frame_finish+0x0/0x1c4 [bridge]
[144516.020026] [<fffffff8041bdbb>] ? br_handle_frame_finish+0x0/0x1c4 [bridge]
[144516.020026] [<fffffff8041bda1>] ? T.865+0x3c/0x56 [bridge]
[144516.020026] [<fffffff8041c107>] ? br_handle_frame+0x188/0x1a1 [bridge]
[144516.020026] [<fffffff8041bf7f>] ? br_handle_frame+0x0/0x1a1 [bridge]
[144516.020026] [<fffffff8126bd81>] ? __netif_receive_skb+0x338/0x4d6
[144516.020026] [<fffffff8126c00b>] ? process_backlog+0xec/0x1c7
[144516.020026] [<fffffff812700a9>] ? net_rx_action+0xa8/0x206
[144516.020026] [<fffffff8104c6ef>] ? __do_softirq+0xc3/0x1a0
```

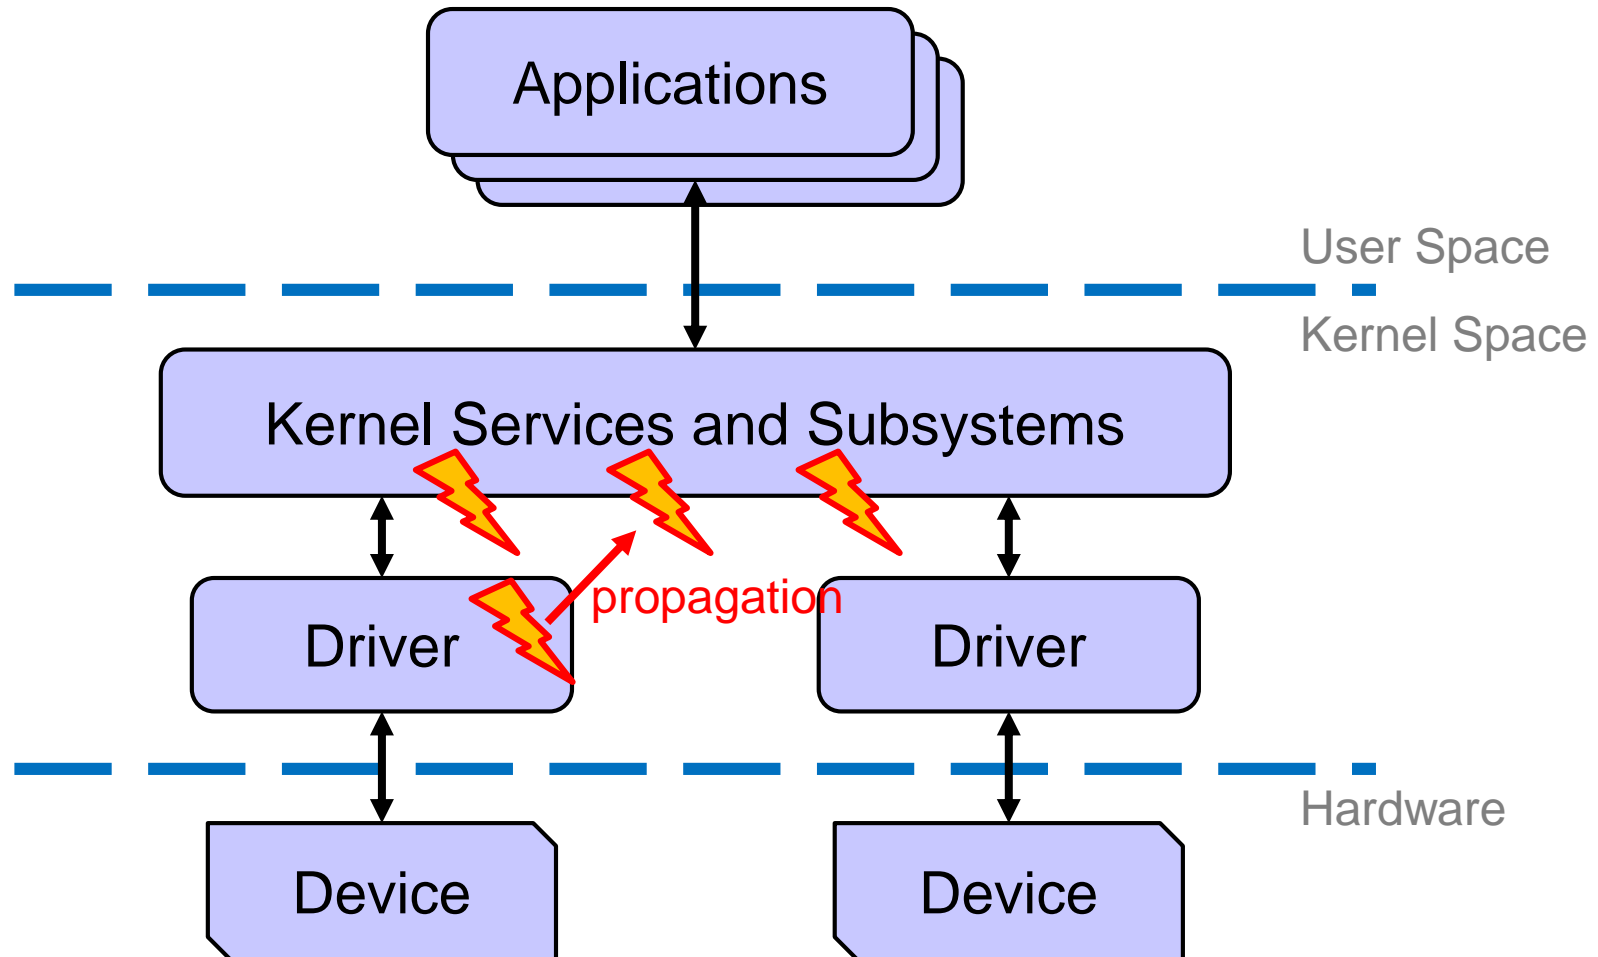
Device Drivers

Major Source of Reliability Issues



Palix et al., "Faults in Linux: Ten Years Later", ASPLOS 2011

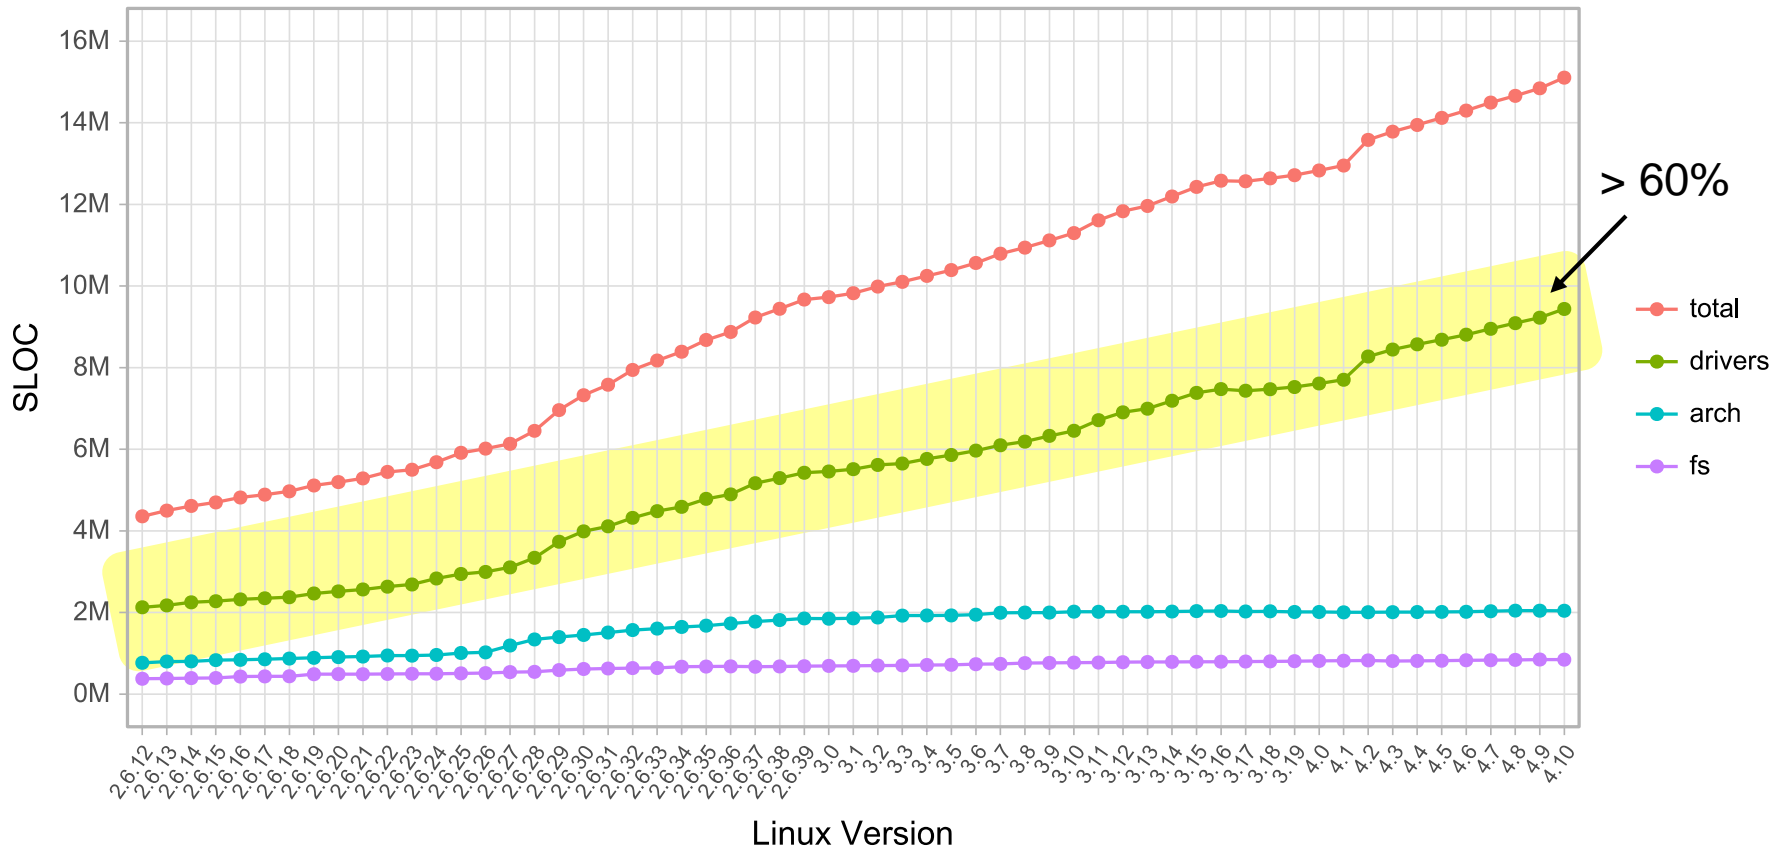
Tight Coupling: Drivers & Kernel



Device Drivers

Major Contributor to Growing Complexity

Lines of Code Development: Linux 2.6.12 to Linux 4.10



generated using 'SLOCCount' by David A. Wheeler

Testing in the Kernel Difficult

It is **compile tested** only because I did not find an easy way **how to run the code**. Well, it should be pretty safe given the nature of the change.

Note, only **compile tested** this as **do not have any hardware** with it in.

I am touching a lot of arch specific code here and I hope I got it right but as a matter of fact I even **didn't compile test** for some archs as I **do not have cross compiler** for them. Patches should be quite trivial to review for stupid compile mistakes though.

— Excerpts from Linux commit messages

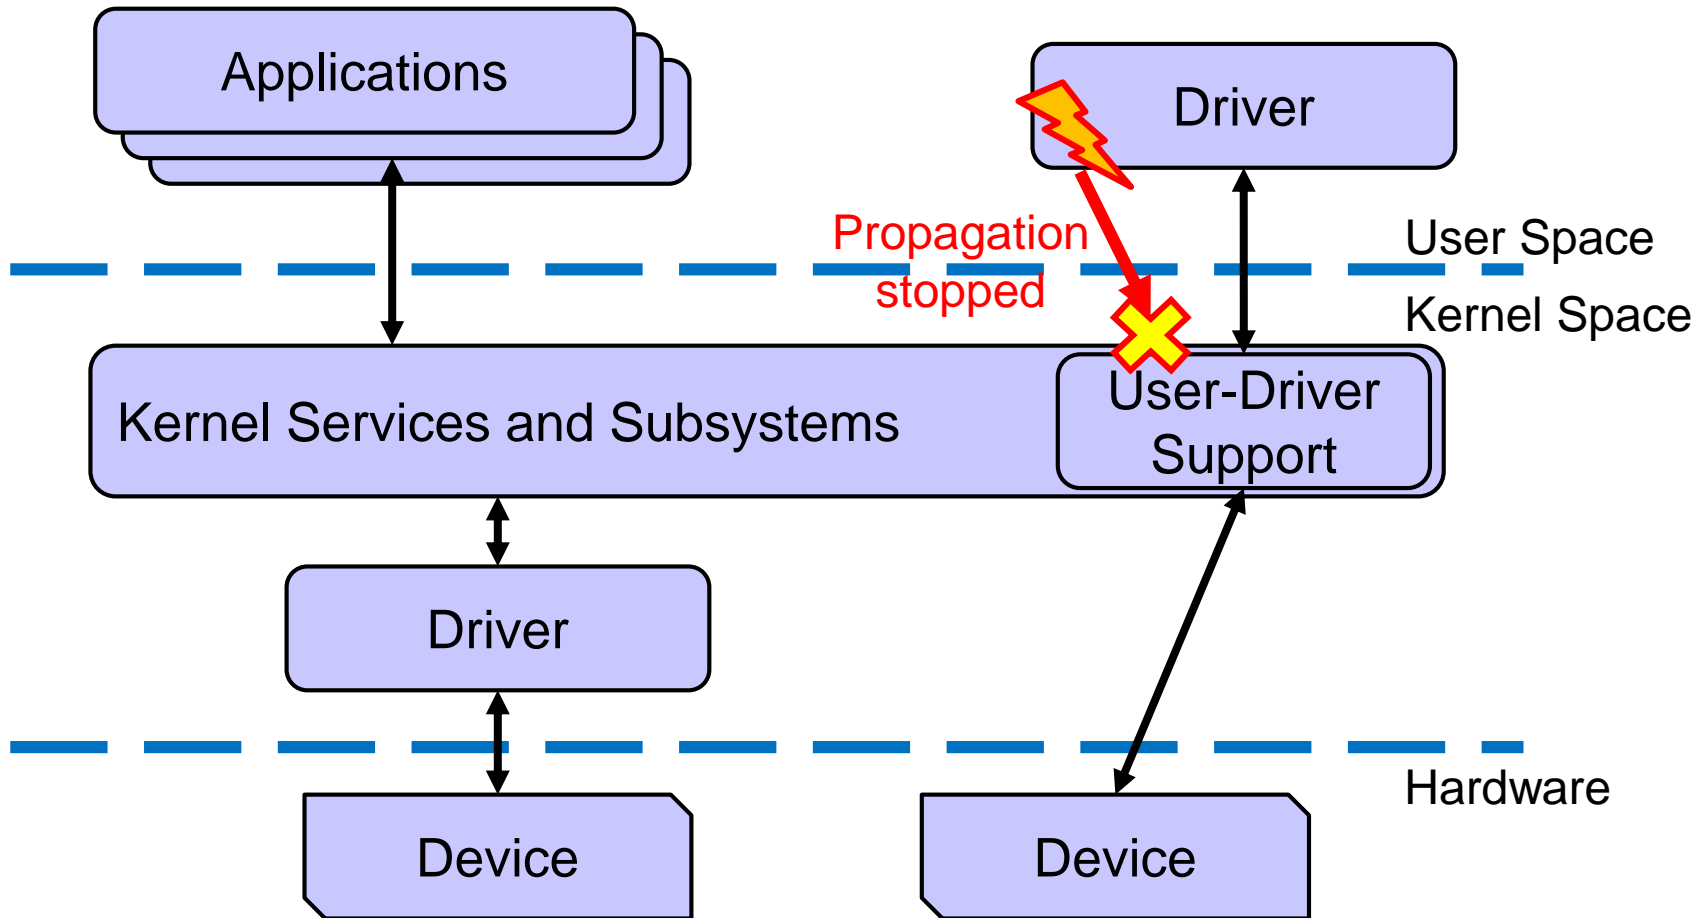
Quick Summary

- Drivers cause reliability issues
- Increasing complexity
- In-kernel development difficult
- Lack of testing



- Focus on 2 directions
 - Tolerate faulty drivers
 - Improve driver code

Tolerating Faulty Drivers: Isolation in User Space



Existing Isolation Approaches

- Several approaches exist
 - None of them has been widely adopted
- ➔ Why is that the case?
- Are they ineffective
 - Too complex
 - Performance impact too high

Are Existing Techniques Ineffective?

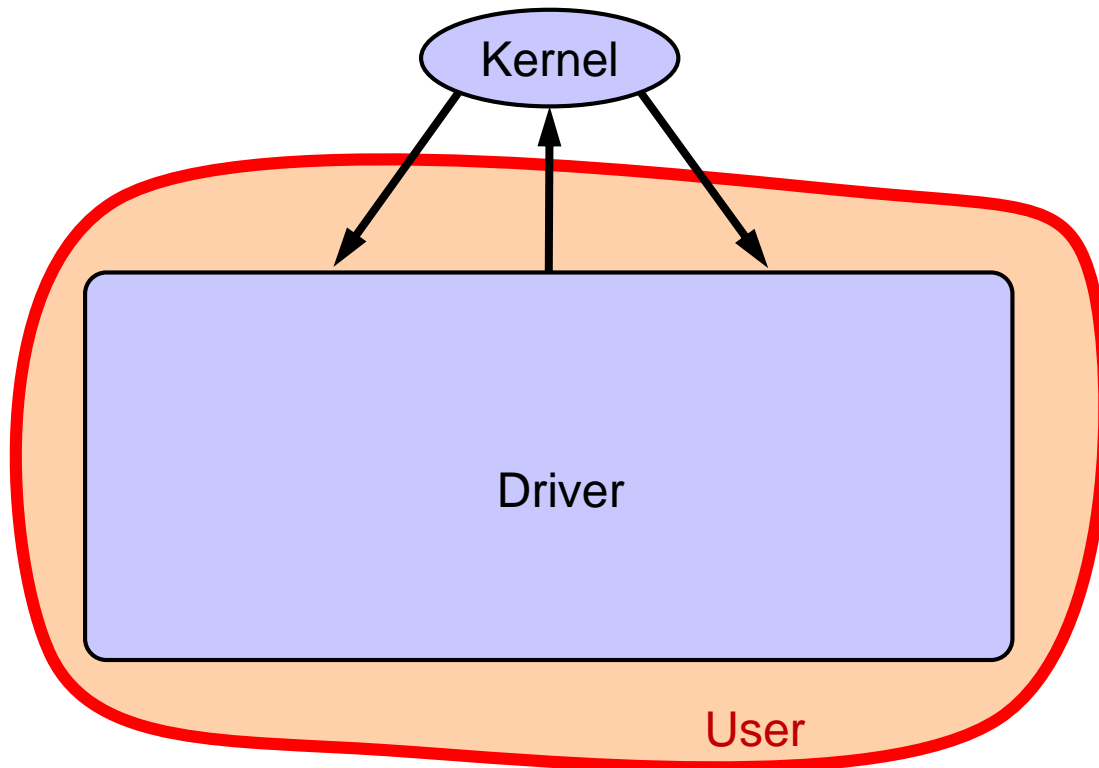
- Unclear how well effects of different fault types are contained
 - Thorough study missing
 - In practice, we may observe surprising effects
 - Which mechanisms are effective against which faults?
 - Protocol violations, hardware faults, memory safety & concurrency bugs...

→ Fault injection study

Isolation Cost Too High?

- Isolation imposes overhead
 - Additional software layer
 - Overhead may vary with different usage scenarios
 - Practical trade-off: Isolate less code, gain performance
- Fine-tune isolation to dynamic usage scenario


Trade-off: Isolation vs Performance




■ Isolation ■ Performance



Trade-off: Isolation vs Performance


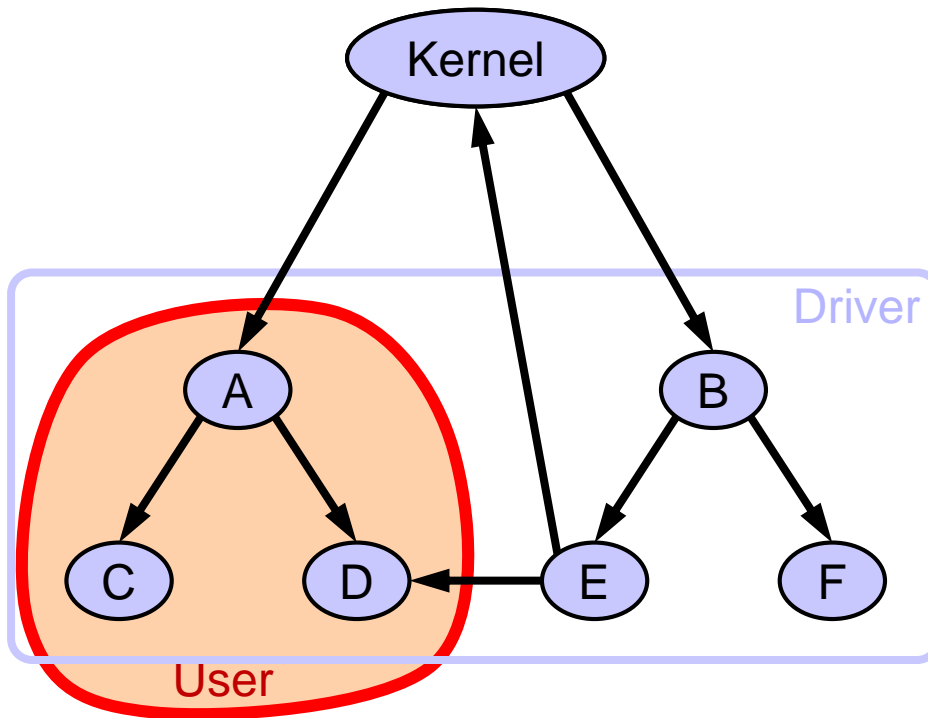


■ Isolation ■ Performance




Trade-off: Isolation vs Performance

■ Isolation ■ Performance



Trade-off: Isolation vs Performance



■ Isolation ■ Performance



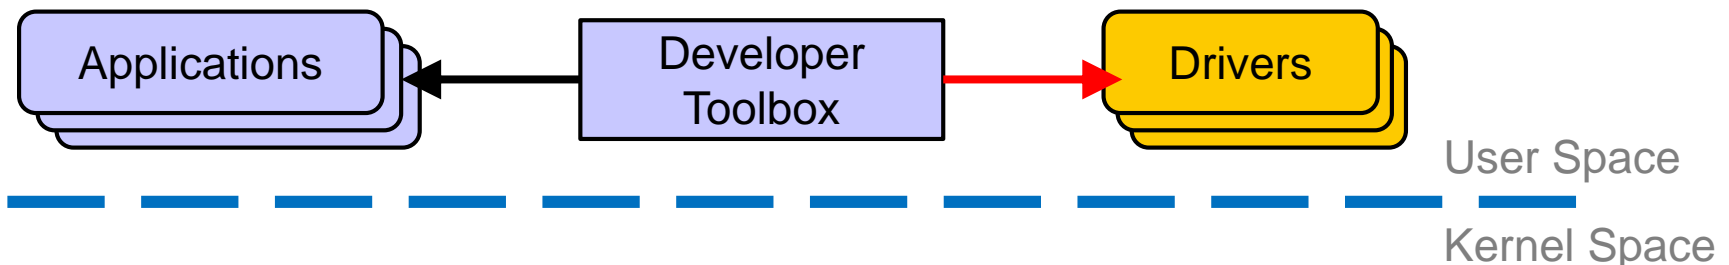
Isolation as Ultimate Answer?

- Isolation contains faults in the driver
 - Driver crashes still annoying
 - Effects on overall system
 - Down times for recovery or restart
 - Decrease functionality, data loss
- ➔ Isolation treats the symptoms, not the root cause

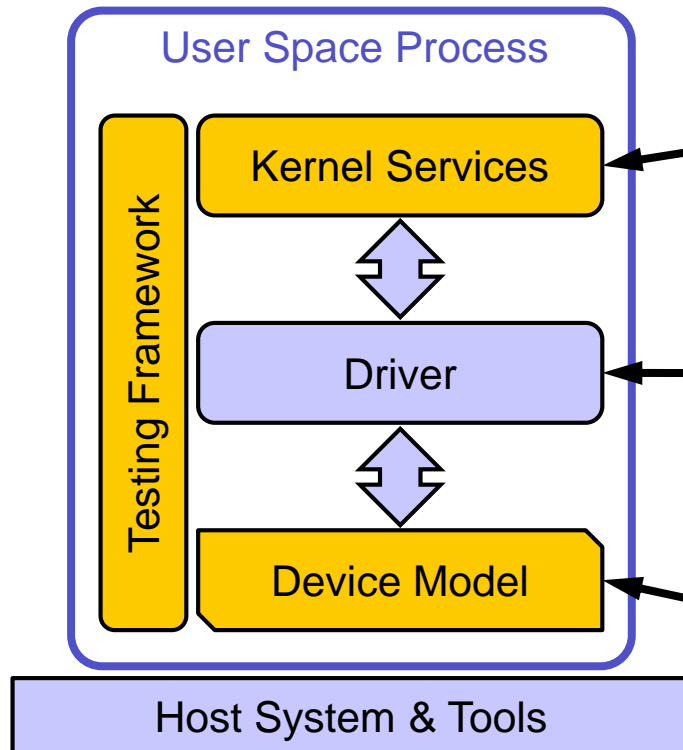
Treating the Root Cause

- We have drivers isolated in user space
- We have a large toolbox for user space development
 - Debugging
 - Testing
 - Code analysis

→ Use the power of the user space toolbox for drivers



Testing in User Space



- Simple Mocking
- Learn from tracing
- User-mode Linux

- Driver “as library”
- Test in isolation
- Test individual functions
- Test for code coverage

- Simple mocking
- Learn from tracing
- Infer from specification

Thank You