

# Initial Evaluation of a User-Level Device Driver Framework

Stefan Götz  
 Karlsruhe University  
 Germany  
 sgoetz@ira.uka.de

Kevin Elphinstone  
 National ICT Australia  
 University of New South Wales  
 kevine@cse.unsw.edu.au

```
*** STOP: 0x0000000A (0x00000000,0x00000002,0x00000000,8038c240)
IRQL_NOT_LESS_OR_EQUAL*** Address 8038c240 has base at 8038c000 - Ntfs.SYS

CPUID: Genuine Intel 6.3.3 irq:1:f SYSVER 0x0000565

Dll Base DateStamp - Name
80100000 336546b1 - ntoskrnl.exe
80001000 33483a53 - atapi.sys
802aa000 33013e6b - epst.spd
802b3000 336015a1 - class.sys
802bd000 33d844ba - Siwvid.sys
f9318000 31ec6c8d - Floppy.SYS
f9488000 31ed868b - RSecDD.SYS
f9319000 3330c02a - i804prt.sys
f947c000 31ec6c94 - kbdcass.sys
f9370000 33248011 - VIDEOPORT.SYS
f9490000 31ec6c6d - vga.sys
f9040000 332480d0 - Ntfs.SYS
a0000000 335157ac - win32k.sys
f60e9000 335bd30a - Fastfat.SYS
fe108000 31ec6c9b - Parallel.SYS
f9050000 332480ab - Serial.SYS


Dll Base DateStamp - Name
80010000 33247f89 - hal.dll
80007000 33248043 - SCSIPTM.SYS
802b5000 336015a2 - Disk.sys
803c0000 335d6537 - Ntfs.sys
803e4000 33d84553 - Ntice.sys
f95e9000 31ec6c99 - Null.SYS
f95ca000 335e00c1 - Beep.SYS
f9474000 332480e4 - auoclass.sys
f95cb000 3379c39d - ctrl2cap.SYS
fe9d7000 3370e7b9 - ati.sys
f93b0000 332480dd - Ntfs.SYS
fe937000 335d6a41 - MDIS.SYS
fe914000 334ea144 - ati.dll
fe110000 31ec7c9b - Parport.SYS
f95b4000 31ec6c9d - ParVdm.SYS

Address dword dump Build [1314] - Name
801afc24 80149905 80149905 ff9eb8c 80129c2c ff9eb94 8025c000 - Ntfs.SYS
801afc2c 80129c2c 80129c2c ff9eb94 00000000 ff9eb94 80100000 - ntoskrnl.exe
801afc34 801240f2 801240f2 ff9edf4 ff9edf0 ff9edc58 80100000 - ntoskrnl.exe
801afc54 80124f16 80124f16 ff9edf0 ff9edc3c 8015ac7e 80100000 - ntoskrnl.exe
801afc64 8015ac7e 8015ac7e ff9edf4 ff9edf0 ff9edc58 80100000 - ntoskrnl.exe
801afc70 80129bda 80129bda 00000000 80088000 8010efc0 80100000 - ntoskrnl.exe

Restart and set the recovery options in the system control panel
or the /CRASHDEBUG system start option. If this message reappears,
contact your system administrator or technical support group.
```


## Why do systems crash?

- Privileged device drivers in monolithic systems
  - No protection
    - system crashes
- [Chou, Engler et al., 01]
  - Many flaws in driver code
  - Large driver code-base
- Run un-trusted drivers?



## Goals

- Improve OS reliability
- Achieve protection
  - Isolate drivers
  - De-privilege drivers
- Application properties
  - user-level drivers
- Maintain performance




## Related Work

- Driver signing
- Isolated but privileged drivers (Nooks)
- “Safe” kernel extensions
- User-level drivers
  - Incomplete bottom-up analysis so far
  - Virtual machines
  - Potential for formal system verification


## Basic Driver Model

- Data transfer & buffer management
- Event notification



## Data Transfer

- Copying expensive → pass-by-reference
- Transfer via shared memory
- Amortization of setup costs




User-Level Device Drivers

7

## Event Notification

- Kernel interaction too expensive
- User-level messaging via shared memory
- Batching




User-Level Device Drivers

8

## Buffer Address Translation

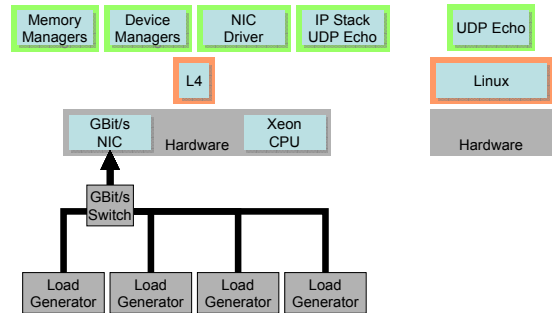
- Physical buffer addresses for DMA
- Translation data unavailable to driver
- Secure state sharing
- Amortization of costs



User-Level Device Drivers

9


## Evaluation: Test Setup



User-Level Device Drivers

10

## Evaluation: Benchmark Results



Similar throughput at lower CPU utilization than Linux

User-Level Device Drivers

11

## Summary

- Poor OS reliability due to device drivers
- Isolated and de-privileged drivers
  - Fault isolation
  - Competitive performance
- Future work
  - Application-level benchmarks
  - Impact of DMA-safe hardware (IO-MMUs)
- *Punch line / Take-home message?*

User-Level Device Drivers

12

## Thank You

- [Chou, Engler et al., 01]: Chou, A., Yang, J., Chelf, B., Hallem, S., Engler, D. *An empirical study of operating systems errors*. In Proceedings of the 18th Symposium on Operating Systems Principles, 2001

User-Level Device Drivers

13

## Session Abstraction

- Interaction and authentication context
- Connects interacting components
- Chainable
- Associates with shared memory
- Relatively long-lived
- Allows for
  - batching
  - granularity vs. performance trade-off

User-Level Device Drivers

14

## Buffer Memory Pinning

- Time-based pinning
  - Pin time guaranteed by memory manager
  - Part of translation cache entries
  - Correct estimates in drivers difficult
- State sharing
  - Pinning requests by drivers
  - Advisory bit in translation cache entries
  - Resource limits enforceable
  - Translation caches writable for drivers

User-Level Device Drivers

15


## Interrupt Handling

- Interrupt delivery to user-level via kernel
- Synchronous IPC as light-weight abstraction
- Overhead from kernel interaction
  - Small compared to off-chip device handling
- Interrupt hold-off techniques
  - Batching in hardware
  - Latency vs. throughput & CPU utilization

User-Level Device Drivers

16

## Execution Time Profile



User-Level Device Drivers

17