What is wrong with Operating Systems?
(and how do we make things better)

FOSDEM 2015

Antti Kantee / Fixup Software Ltd.
pooka@rumpkernel.org
@anttikantee

http://cm.bell-labs.com/cm/cs/who/dmr/picture.html

From computers to computing

“artist's” rendition of all software
(actual implementation may vary)

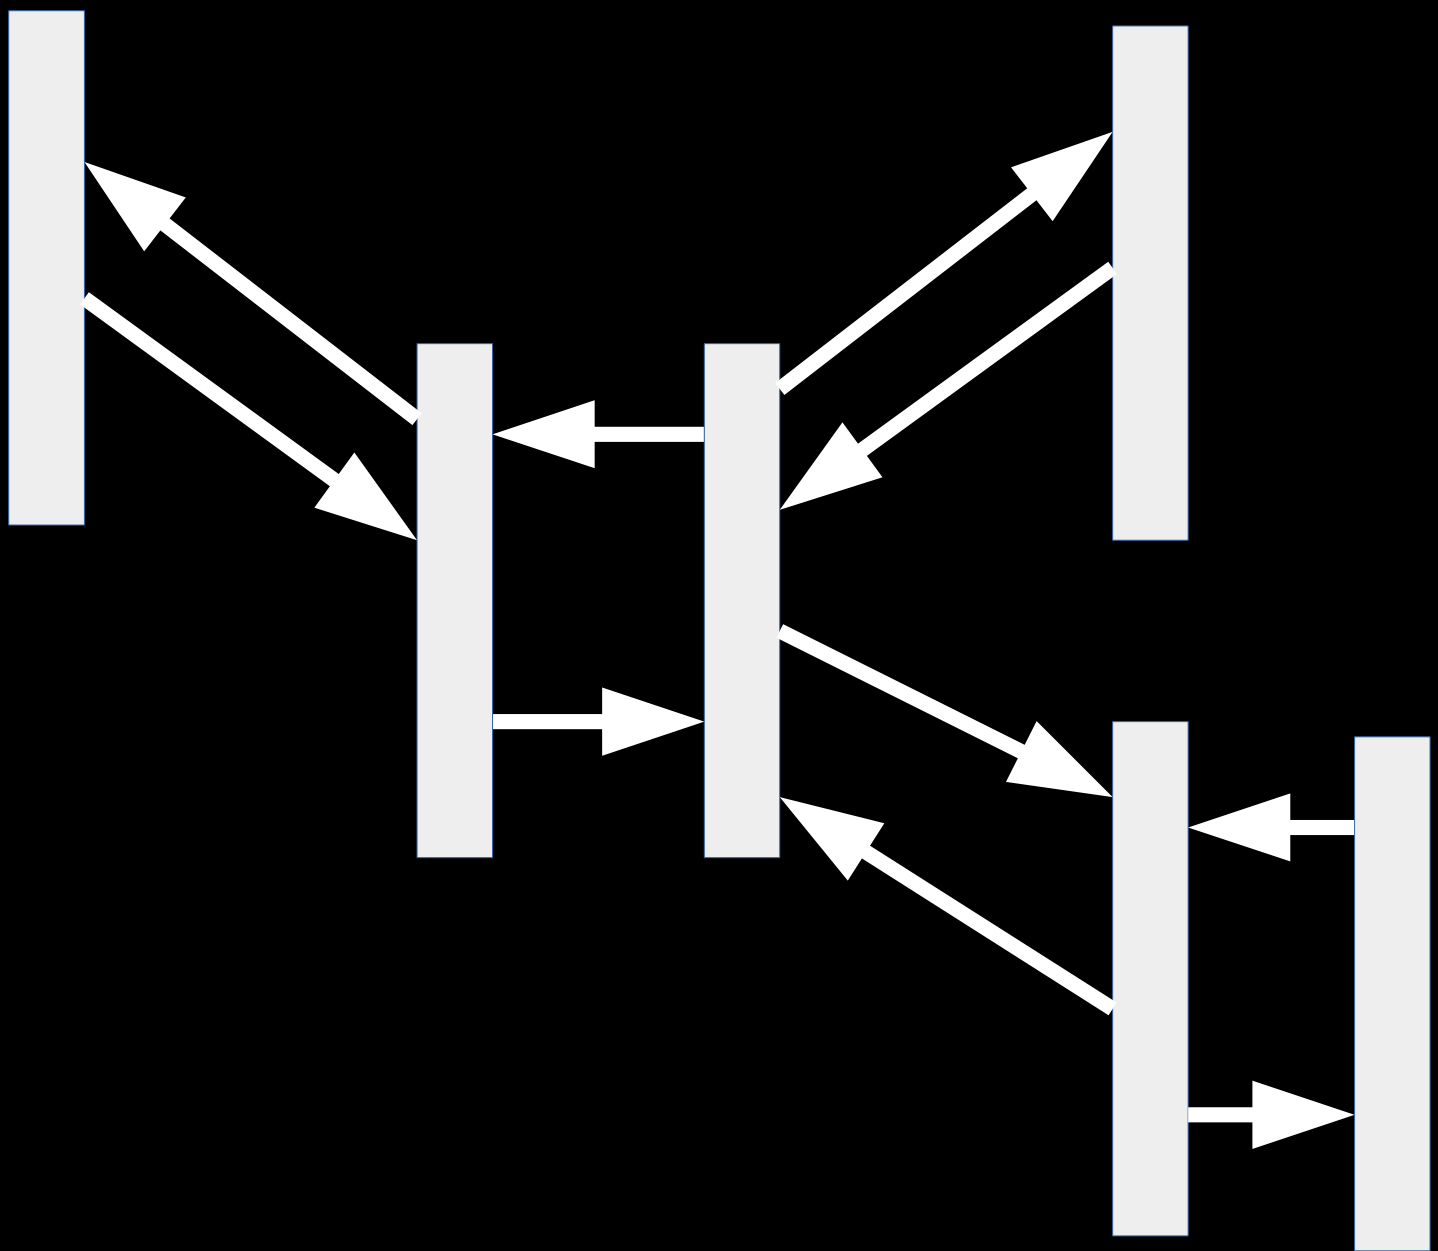

What does an OS do?

Supporting applications

1) enabling
2) abstracting
3) isolating
4) monitoring

+ orchestrating

What does an OS do?

Supporting applications

1) enabling => getting ready
2) abstracting => drivers
3) isolating => can't do
4) monitoring => info

\+ orchestrating

I have no quarrel with you, good sir OS, but I must cross this bridge

Monty Python

Monty Python

Monty Python

Handle normal and worst cases separately

The normal case must be fast.

The worst case must make some progress.

Butler Lampson
Hints for Computer System Design (1983)

Handle normal and worst cases separately

In most systems it is all right [...] even to deadlock the entire system, as long as this event is detected automatically and doesn’t happen too often. The usual recovery is by crashing some processes, or even the entire system. At first this sounds terrible, but one crash a week is usually a cheap price to pay for 20% better performance. Of course the system must have decent error recovery, but that is required in any case, since there are so many other possible causes of a crash.

Butler Lampson
Hints for Computer System Design (1983)

A good driver is hard to write

The need for drivers does not imply the need for an OS

Every problem in operating systems can be solved by *removing* layers of indirection

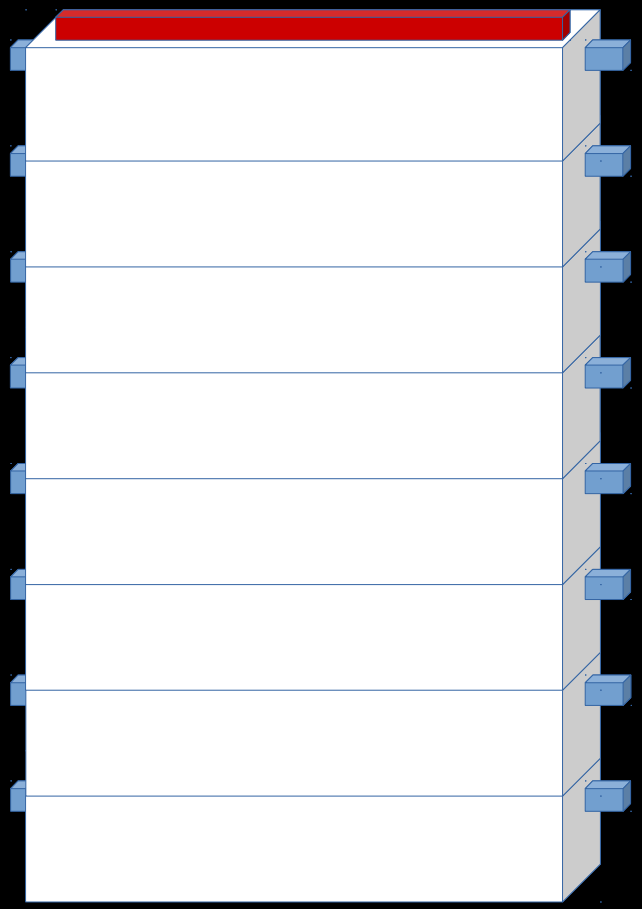

"artist's" rendition of poorly architected lasagna

“There's no such thing as #1”

Userspace
vs
Kernelspace

Microkernels
vs
Virtualization

A good driver is hard to write
is easy to run

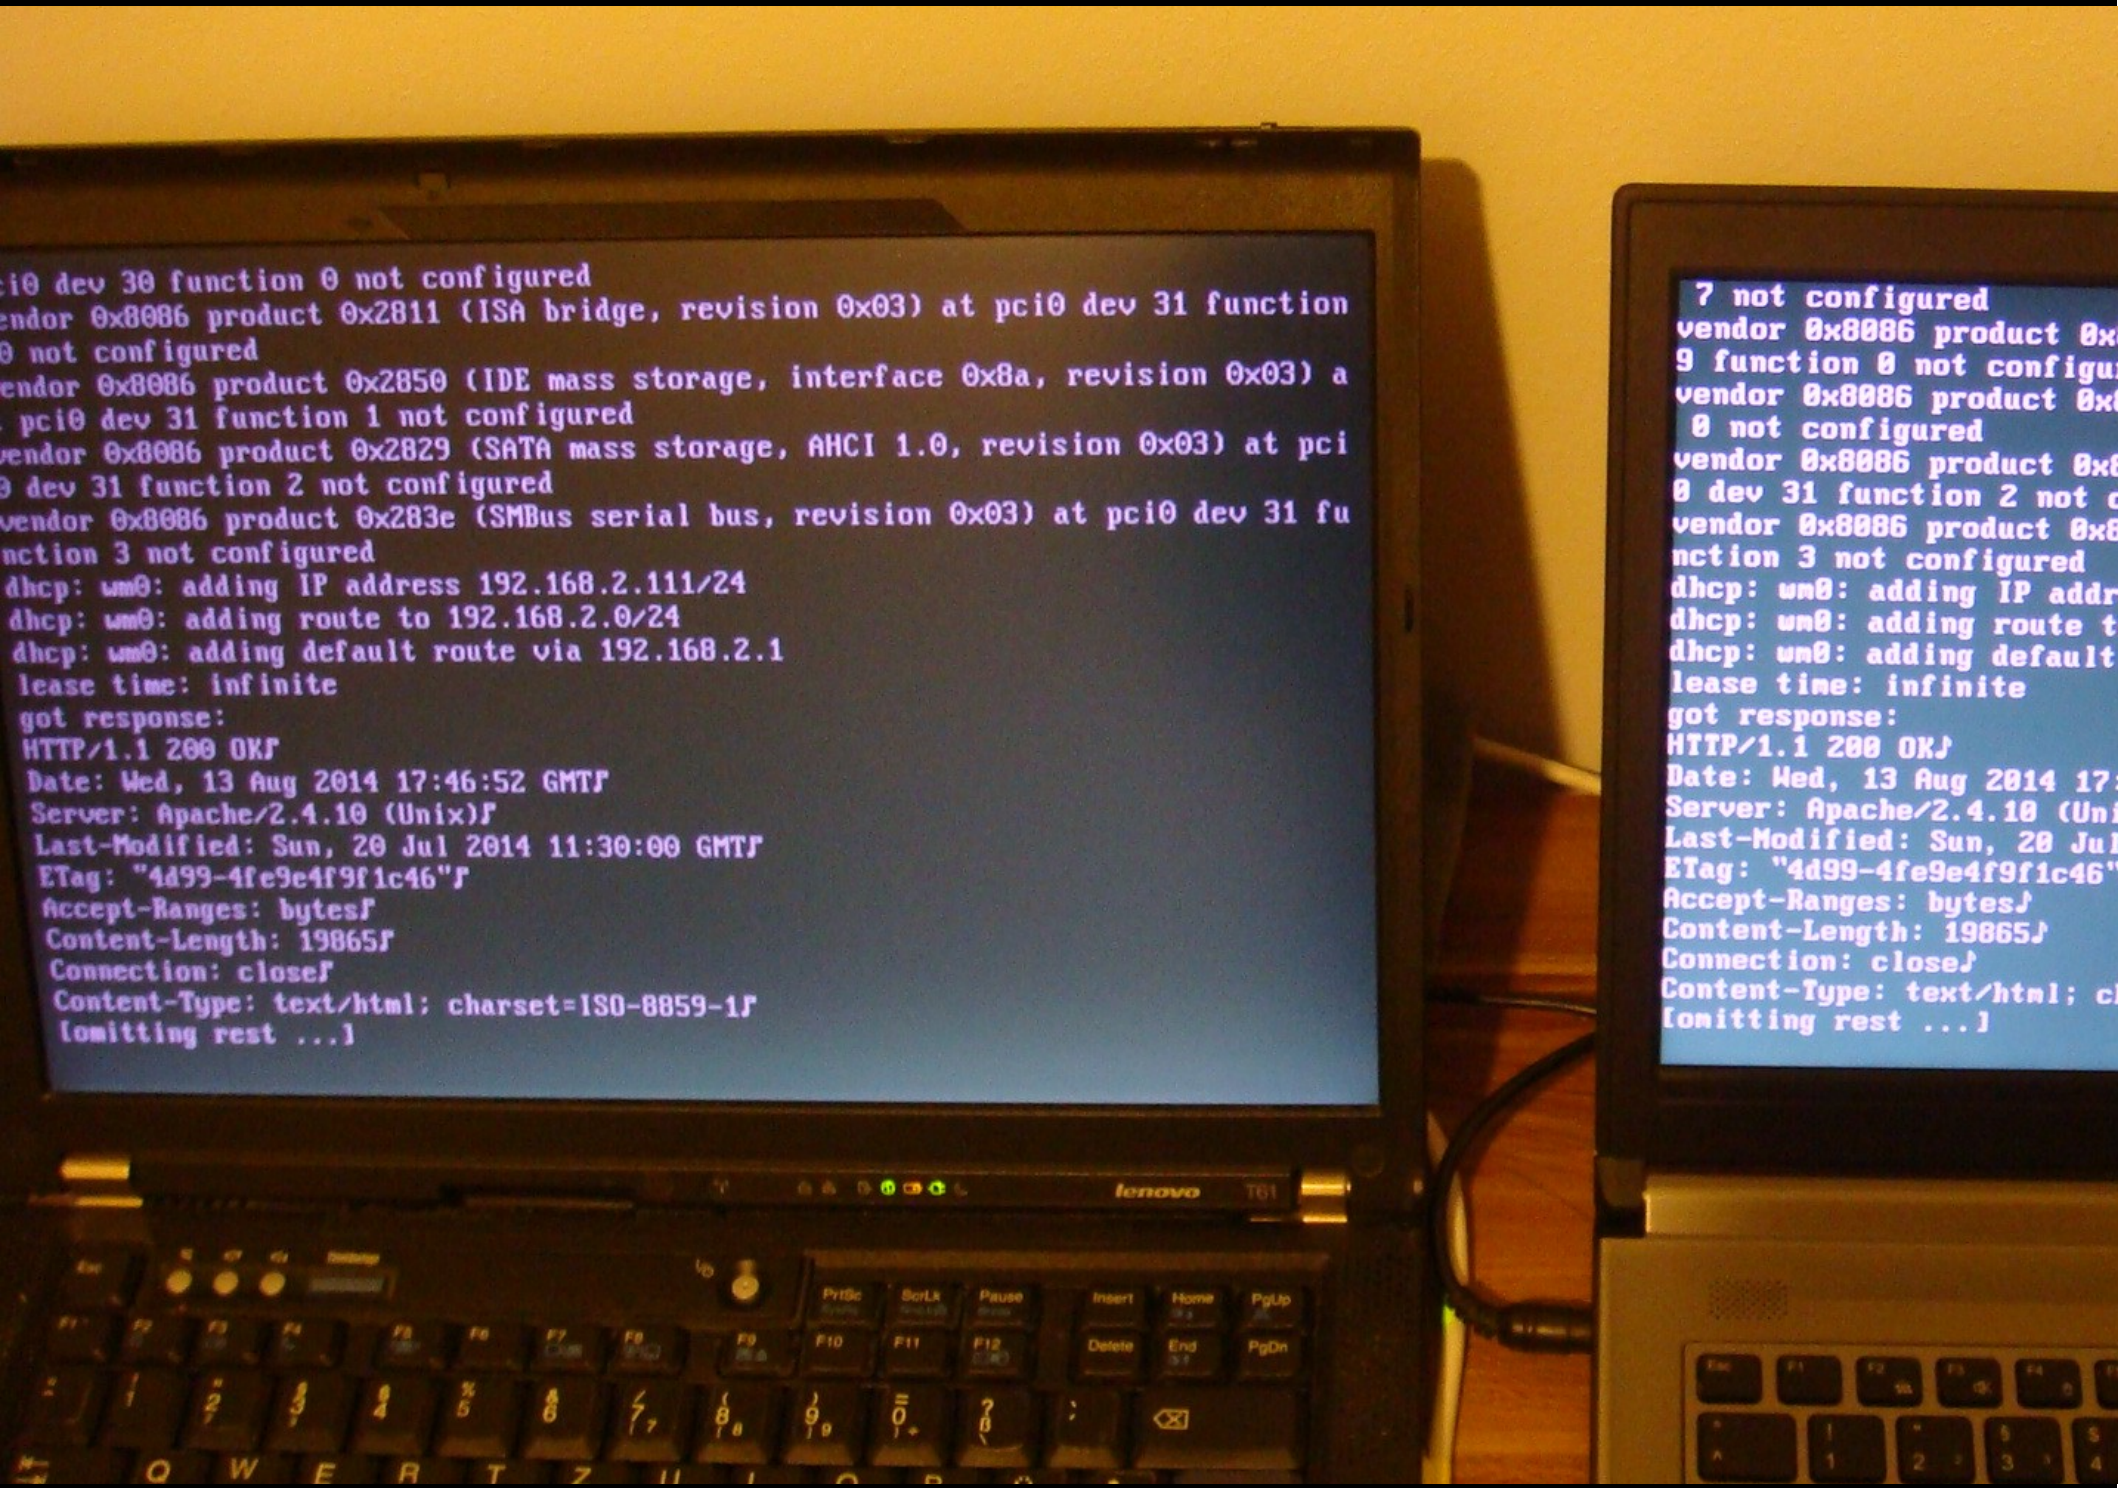
i0 dev 30 function 0 not configured
endor 0x8086 product 0x2811 (ISA bridge, revision 0x03) at pci0 dev 31 function
0 not configured
endor 0x8086 product 0x2850 (IDE mass storage, interface 0x8a, revision 0x03) a
pci0 dev 31 function 1 not configured
endor 0x8086 product 0x2829 (SATA mass storage, AHCI 1.0, revision 0x03) at pci
0 dev 31 function 2 not configured
endor 0x8086 product 0x283e (SMBus serial bus, revision 0x03) at pci0 dev 31 fu
nction 3 not configured
dhcp: wm0: adding IP address 192.168.2.111/24
dhcp: wm0: adding route to 192.168.2.0/24
dhcp: wm0: adding default route via 192.168.2.1
lease time: infinite
got response:
HTTP/1.1 200 OK
Date: Wed, 13 Aug 2014 17:46:52 GMT
Server: Apache/2.4.10 (Unix)
Last-Modified: Sun, 20 Jul 2014 11:30:00 GMT
ETag: "4d99-4fe9e4f9f1c46"
Accept-Ranges: bytes
Content-Length: 19865
Connection: close
Content-Type: text/html; charset=ISO-8859-1
[omitting rest ...]
lenovo
T61
7 not configured
vendor 0x8086 product 0x
9 function 0 not configu
vendor 0x8086 product 0x
0 not configured
vendor 0x8086 product 0x
0 dev 31 function 2 not
vendor 0x8086 product 0x
nction 3 not configured
dhcp: wm0: adding IP addr
dhcp: wm0: adding route t
dhcp: wm0: adding default
lease time: infinite
got response:
HTTP/1.1 200 OK
Date: Wed, 13 Aug 2014 17
Server: Apache/2.4.10 (Un
Last-Modified: Sun, 20 Ju
ETag: "4d99-4fe9e4f9f1c46"
Accept-Ranges: bytes
Content-Length: 19865
Connection: close
Content-Type: text/html; c
[omitting rest ...]

Hardware

(software would be so much fun without it)

Recap

1: application('s) rules
2: when in doubt, leave it out
3: bad granularity
4: bootstrap and evolve
5: if all you have is a hammer,
you'll also have a sore thumb

?