BPF
All your Packets belong to Me

@_xhr_

xhr@giessen.ccc.de

BPF ?

tcpdump ?

Packet Flow

NIC

↓

Link-Layer Driver

↓

Protocol Stack

↓

Userland

Smart Idea

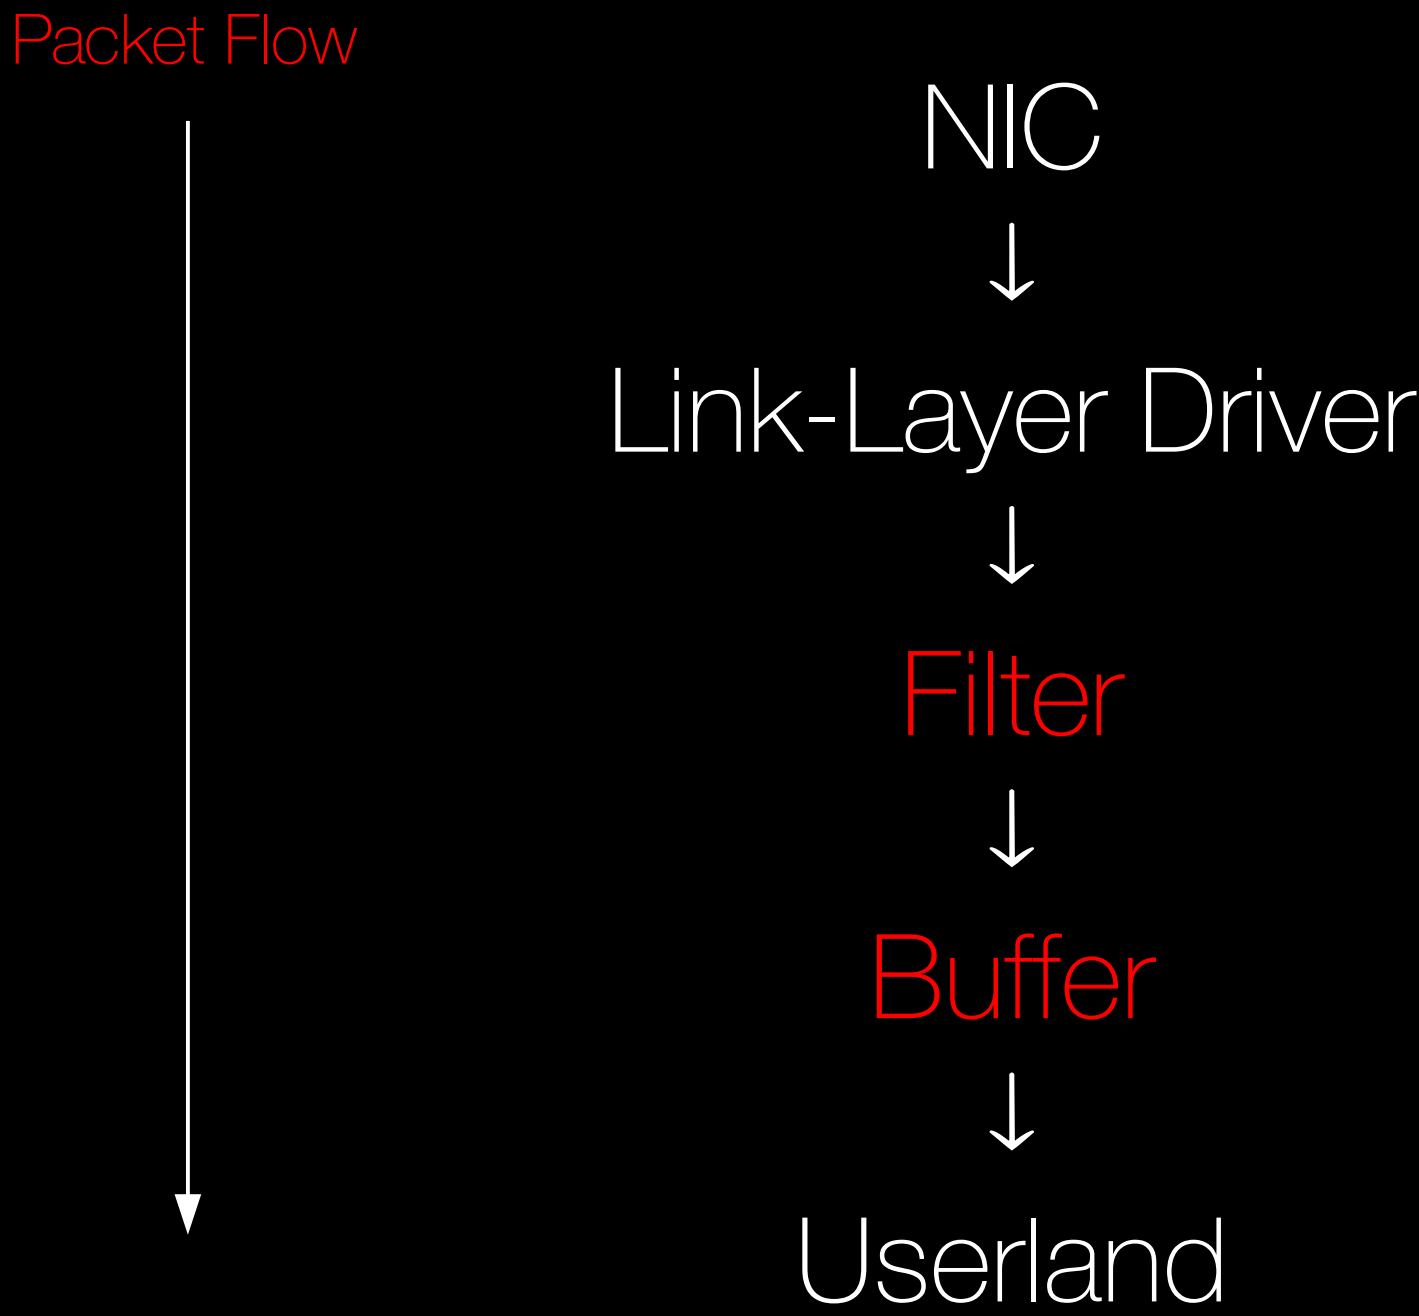
Packet Flow
NIC
↓
Link-Layer Driver
↓
Filter
↓
Buffer
↓
Userland

BPF is rather old...

McCanne. Jacobson. *The BSD Packet Filter: A New Architecture for User-level Packet Capture.* in USENIX. 1993.

Have you met ...

```
tcpdump -i eth0 ip6
```

That's the filter

Ethernet Protocol Type
0x86dd == IPv6
0 ldh [12]
1 jeq #0x86dd jt 2 jf 3
2 ret #65535
3 ret #0
Accept Packet
Drop Packet

Linux got a BPF JIT in 2011

Check `net/core/filter.c`

Packet Filter only for Packets?

seccomp?

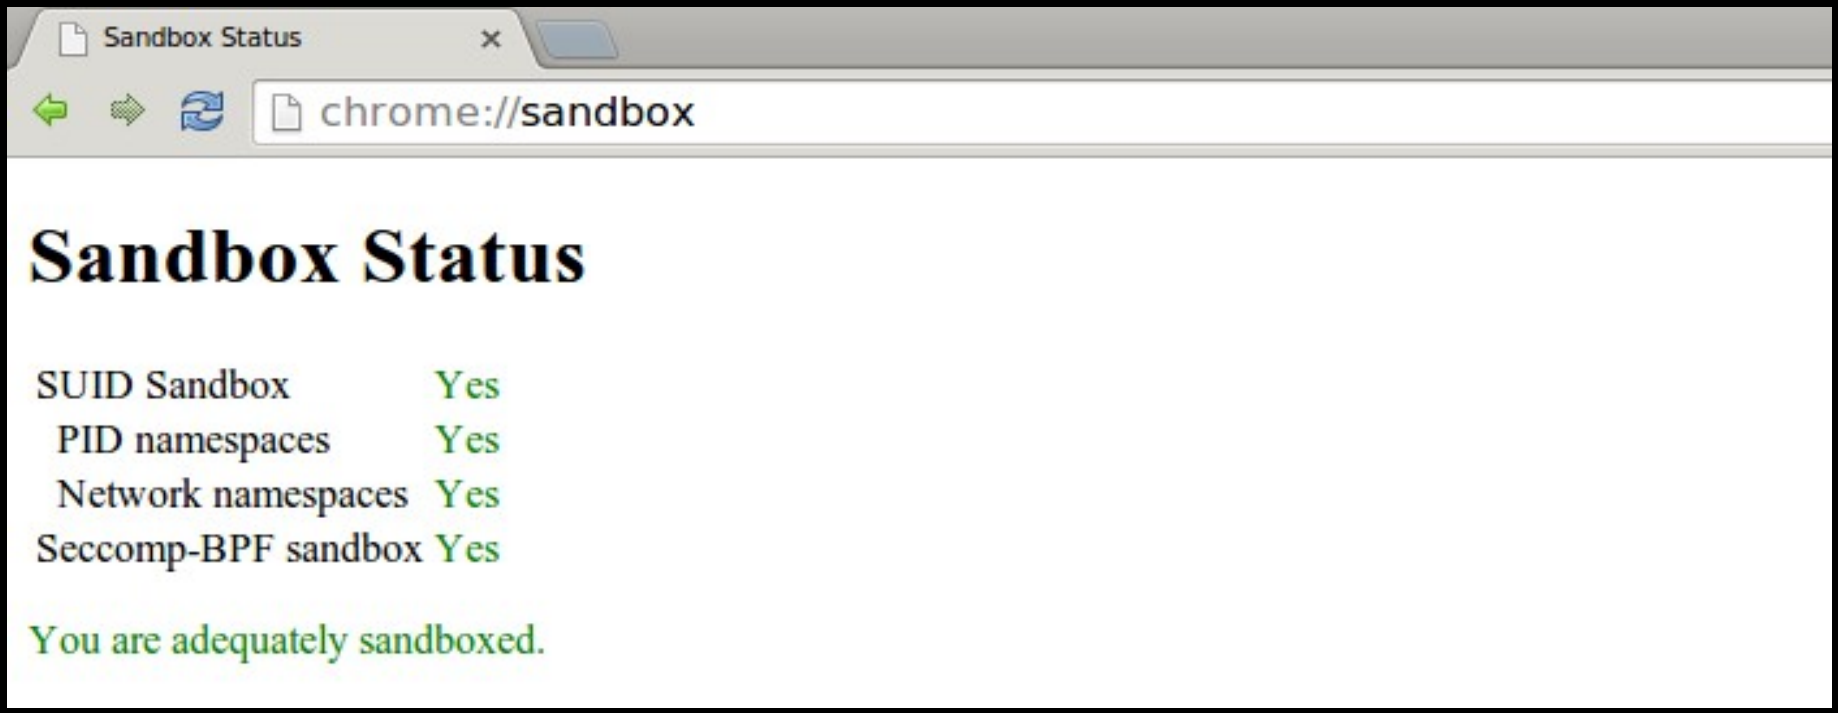
Sandbox Status
chrome://sandbox
Sandbox Status
SUID Sandbox Yes
PID namespaces Yes
Network namespaces Yes
Seccomp-BPF sandbox Yes
You are adequately sandboxed.

So, how does this work?

Attach a filter to a socket

```
[...]

struct sock_filter code[] = {
        { 0x28,  0,  0, 0x0000000c },
          [...]
};

struct sock_fprog bpf = {
   .len = ARRAY_SIZE(code),
   .filter = code,
};

sock = socket(PF_PACKET, SOCK_RAW,
   htons(ETH_P_ALL));

ret = setsockopt(sock, SOL_SOCKET,
   SO_ATTACH_FILTER, &bpf, sizeof(bpf));

[...]
```

So, how can I use this?

Need for Space

A	32 bit wide accumulator
X	32 bit wide X register
M[]	16 x 32 bit "scratch memory"

Some Instructions

`ld*`	Load Instructions
`st*`	Store Instructions
`j*`	Jumps
`ret`	Return
`$alu`	ALU instructions

Hmm … k. IDE anyone?

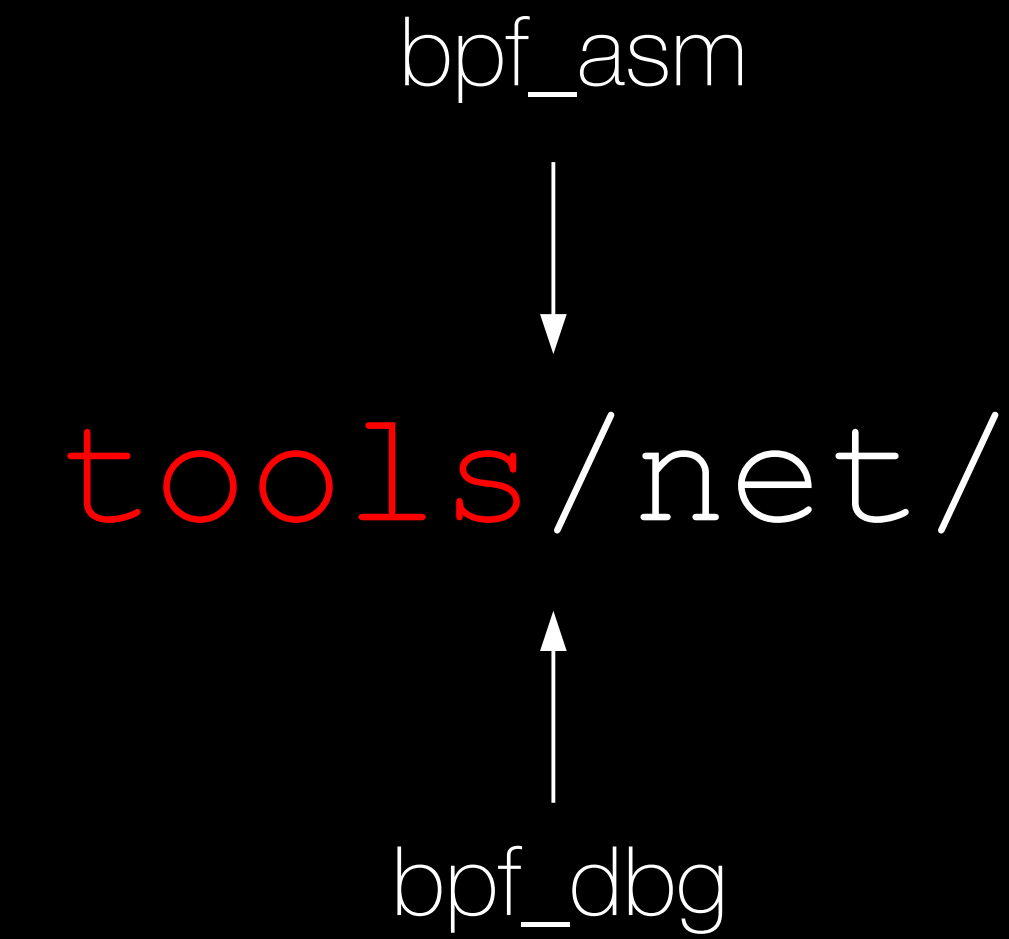
bpf_asm
tools/net/
bpf_dbg

Demo

What now?

Packet Filtering

Can I haz xt_bpf, plz?

```
iptables -A <CHAIN> \
  -m bpf \
  --bytecode "…" \
  -j <TARGET>
```

Because we can!!1

Full packet control

And Why?

Fine grained filters

Q & A

xhr

xhr@giessen.ccc.de
@_xhr_