sockdev - socket to net device redirection

Linux Kernel Lab
Winter Term 2007/2008

Elmar Hoffmann <elho@elho.net>
Martin Henze <martin.henze@rwth-aachen.de>

RWTH Aachen University

March 28, 2008

Overview

Motivation & Goals

How it works

Live demo

Future Work

Questions

Motivation & Goals

- ethernet tunnel
  - full-blown device supporting arbitrary protocols
    - virtualization
    - network simulation
  - ability to be bridged with other ethernet devices
- kernel space
  - better performance
    - avoid unnecessary copying of packets to user space and back

How it works

- user space process creates and connects or binds socket
- user space process passes control over the socket to sockdev
  - currently using `ioctl(2)` on `/dev/sockdev`
- sockdev passes packets arriving on the socket to the net device and vice versa in kernel space
  - `net_device->hard_start_xmit()`
  - `sock->sk_data_ready()`

Live demo

- VM network: `192.168.150.0/24`
- sockdev network: `192.168.0.0/24`

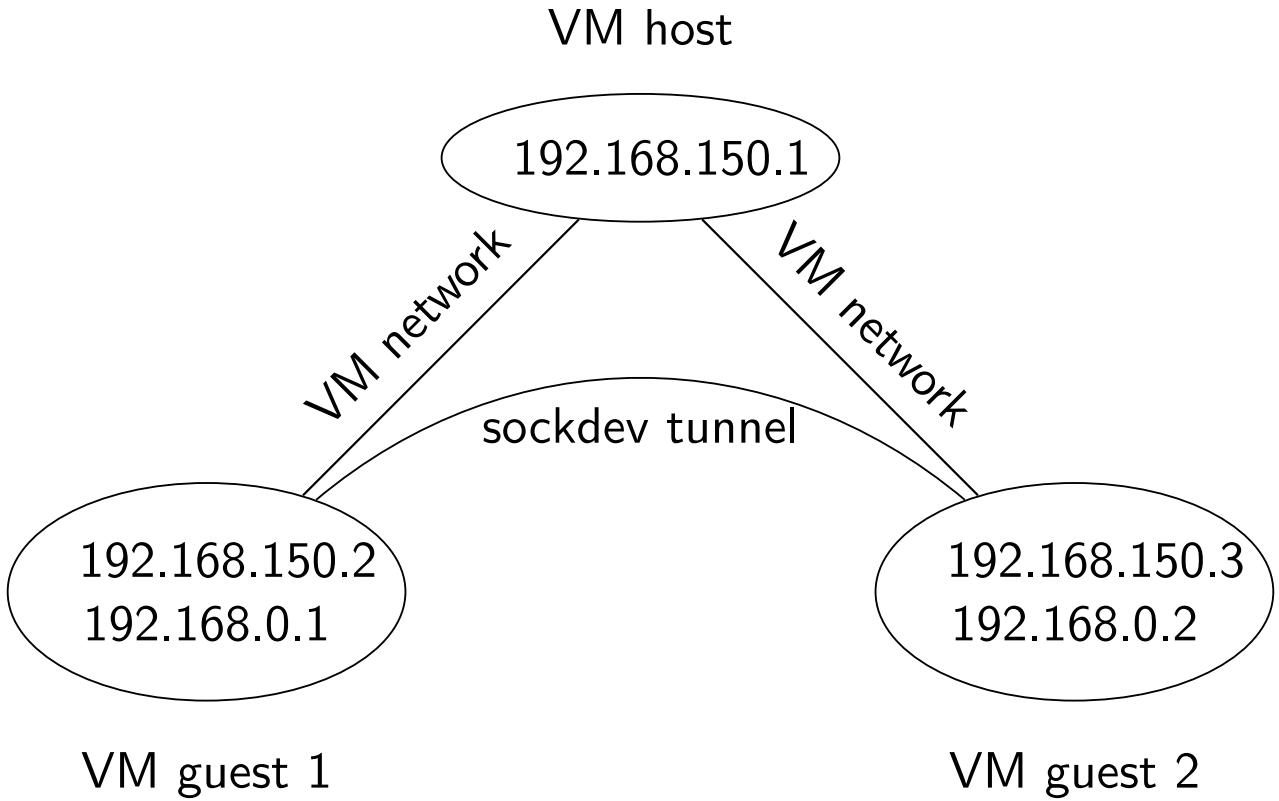

Future Work

- code cleanup
- interface statistics
- support standard encapsulations
  - GRE
  - EtherIP
- make it a kernel patch
  - use `setsockopt(2)` instead of `ioctl(2)`
  - use own private field to store data instead of `sock->sk_user_data`

Resources

- "sockdev - socket to net device redirection" slides

  `http://www.elho.net/pub/`

- sockdev source code

  `http://www.elho.net/dev/lkl/`

Questions?