



# Breaking Open Linux Switching Drivers

**Andy Gospodarek**

---

Cumulus Networks

16 Feb 2015

netdev01

---

# Agenda

- Why
- History
- Design proposal
- Future possibilities

---

# Why am I here?

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Linux kernel should enable others to  
create the next generation of forwarding  
devices**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# **Integrate support for offload hardware directly into the the Linux kernel**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# Hardware Platform History

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Market dominated by switch and  
router vendors providing expensive  
proprietary solutions**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Proprietary software running on switches  
and routers was not open for developers  
and users to enhance**

---

**Today's hardware platforms are  
significantly higher-performance and more  
generally available**

---

# **Spare CPU cycles are available for applications to run directly on the switch**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Bare-metal platforms are now appealing  
and available to commercial Linux  
vendors, developers, and users**

---

# Software History

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# **15+ years with Linux as a viable OS for host processor on switches/routers**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# 10+ years Linux “support” by ASIC vendors

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# Basic in-kernel switching/offload layer support in v3.19

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Software architecture to control ASICs  
has not fundamentally  
changed in the last decade**

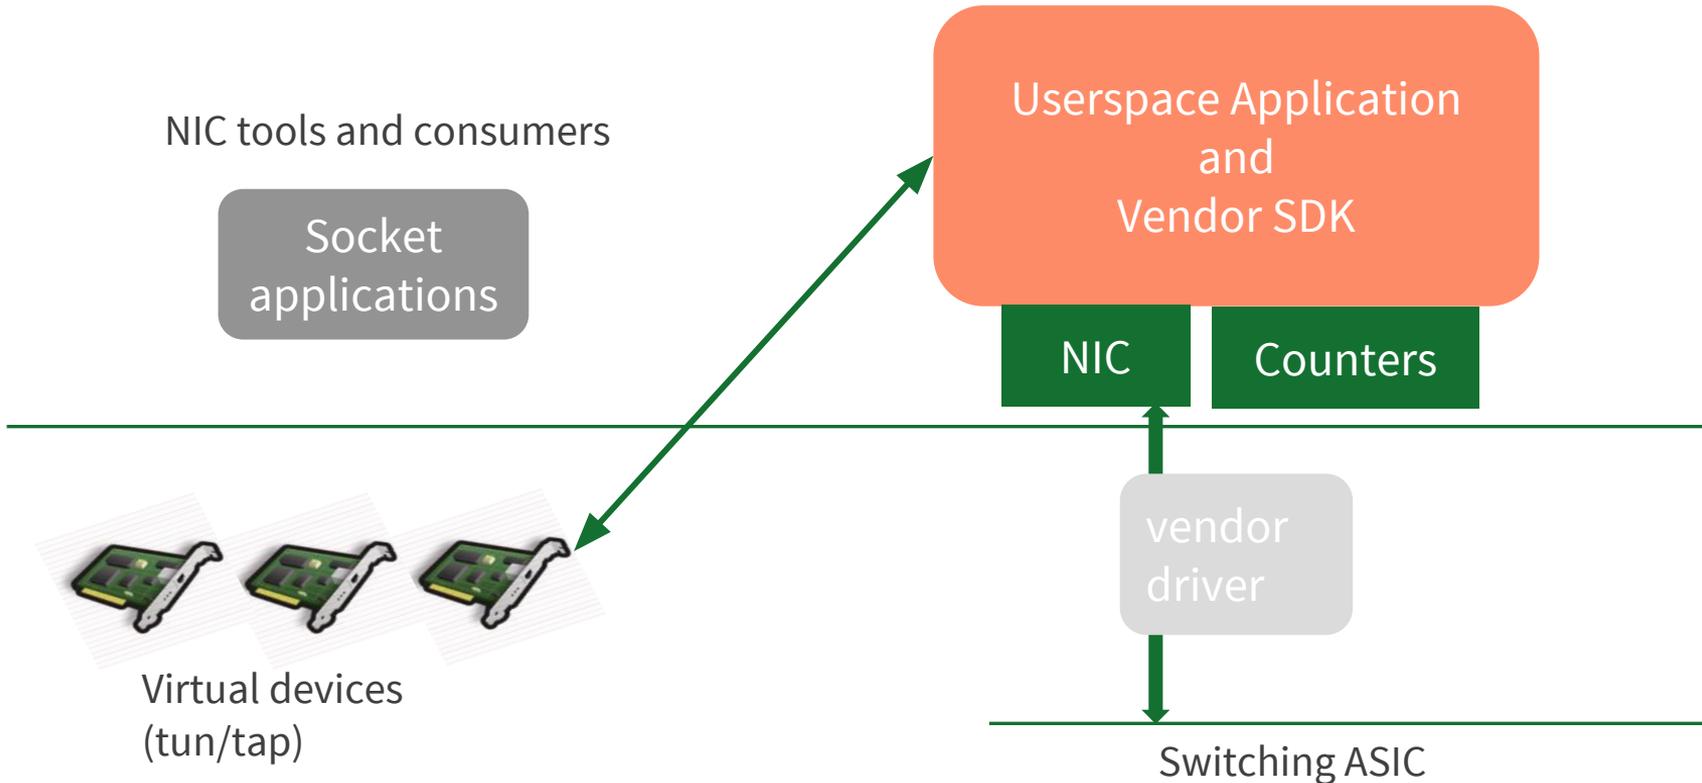
*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# What exactly does that look like?

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

# Typical packet path



*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

Routing suites, bridge controllers, etc

iproute2 and friends

Userspace Application and Vendor SDK

FDB/FIB    Init

Linux kernel structures



Virtual devices (tun/tap)

vendor driver

Switching ASIC

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Painful for those developing switches as  
management applications need to talk to  
kernel/netlink and SDKs**

Routing suites, bridge controllers, etc



iproute2 and friends

Userspace Application and Vendor SDK

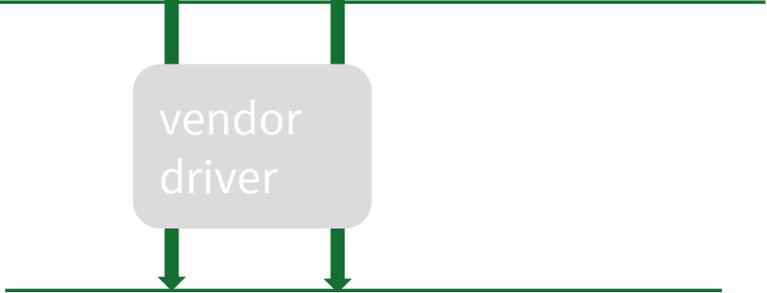
FDB/FIB    Init

Linux kernel structures



Virtual devices (tun/tap)

vendor driver



Switching ASIC

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Much better design, but each SDK  
supported needs a new translation  
between netlink and SDK**

---

# **Kernel hackers and distribution vendors see a simple solution**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# Get rid of all closed-source SDKs

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Great idea!**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# **Vendors do not want to open source their SDKs**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# Can we use a userspace SDK and a kernel driver at the same time?

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*



*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Not if you want it upstream!**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**OK...how do I get started?**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

## Phased Approach

- Participate!
- Pick a hardware platform
- Write and post a switchdev-compatible network driver
- Enhance that driver to add `ndo/offload_ops` to driver

---

# **Attend conferences, participate on mailing-lists, and post patches**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

# Write and post a switchdev-compatible network driver

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

## Advantages

- Provide network access via front panel ports
- Phased approach to working upstream
- Applications can be developed without need for hardware offload

---

# What might that look like?

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

NIC tools and consumers

Socket  
applications

Ethtool

Port Init

Link Mgmt

RxTx

Counters



switchdev compatible driver

Switching ASIC

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**Great, we are upstream!**  
**Are we done?**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**No!**

---

# **Add offload support to driver as upstream infrastructure is developed**

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

What might *that* look like?

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

NIC tools, Routing suites, bridge controllers, etc

Socket applications

Ethtool

iproute2 and friends

Port Init

Link Mgmt

Counters

RxTx

Offload Ops



switchdev compatible driver

Switching ASIC

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**“If you are the first you will be so cool.”**

**-DaveM**

---

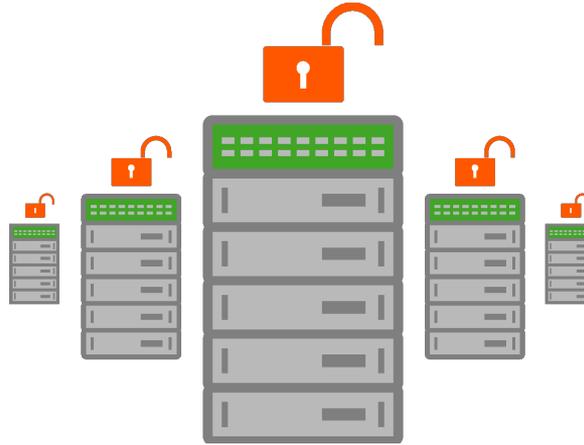
# Get coding

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*

---

**“..and we’ll help you maintain it”**

**-DaveM**



# Thank You!

---

© 2013 Cumulus Networks. Cumulus Networks, the Cumulus Networks Logo, and Cumulus Linux are trademarks or registered trademarks of Cumulus Networks, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners. The registered trademark Linux® is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis.

*Proceedings of netdev 0.1, Feb 14-17, 2015, Ottawa, On, Canada*