



Linux Networking for the Enterprise

How linux REALLY runs your entire enterprise

April 6th, 2017

Shrijeet Mukherjee | Cumulus Networks



Datacenter and Linux





Linux boundaries today

APP

- Hadoop/Zookeeper

Hypervisor

- KVM/Xen

Containers


- Docker/Mesos/Kubernetes

Host

- Bridges/Routers/OVS

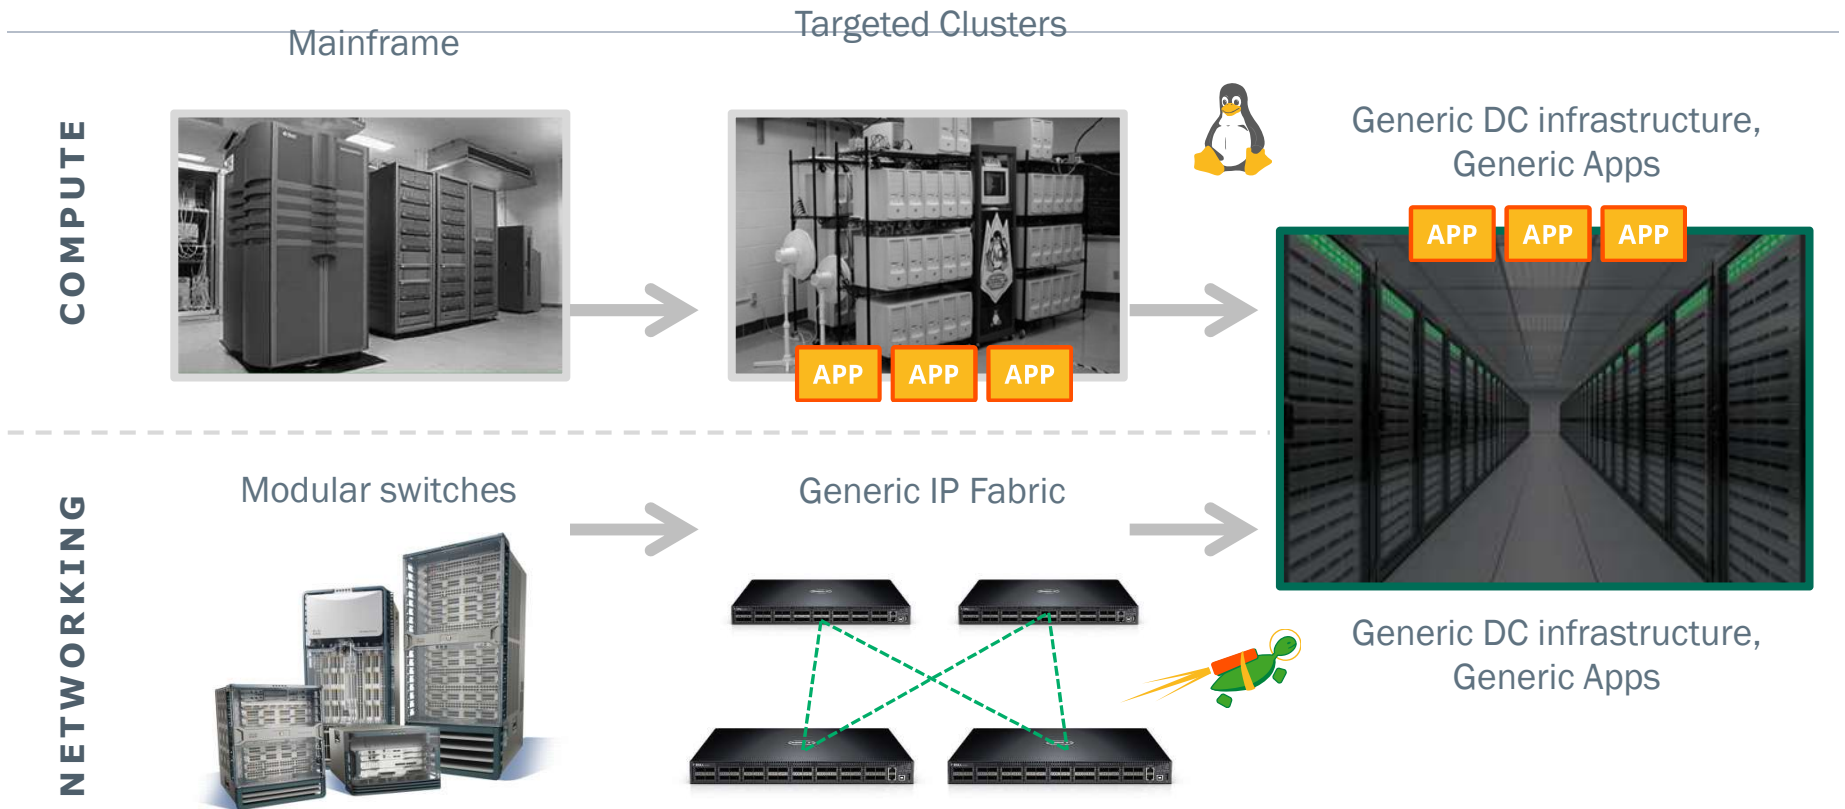
Each builds on each other and is silently interchangeable i.e Infra

- Even MSFT picked Linux for it's cloud OS






Time Travel to the Modern Data Center



 CUMULUS



Get Linux in here !!



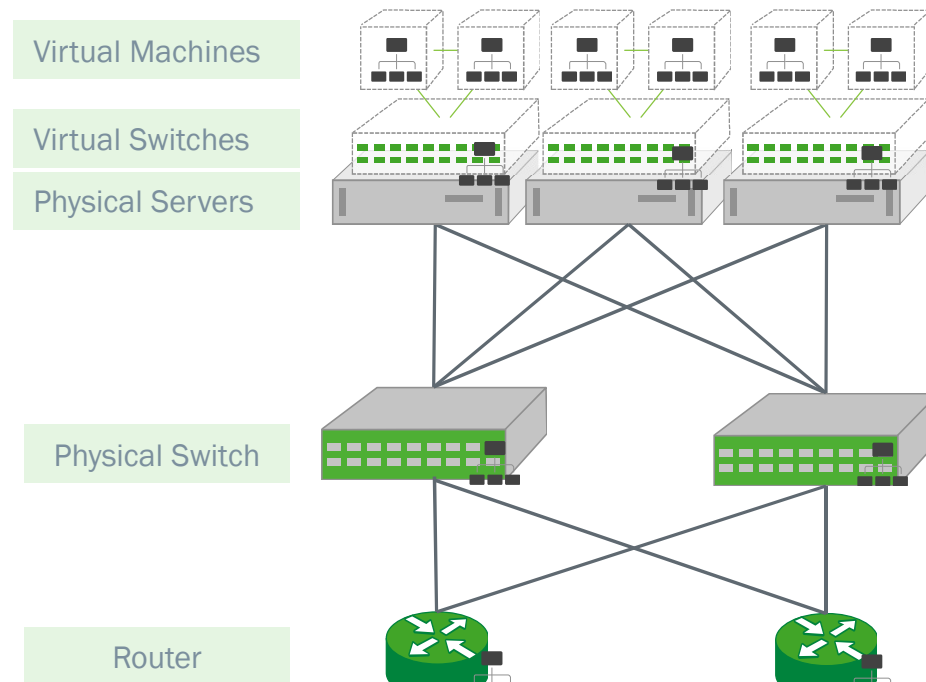


Linux boundaries today

Physical Switches and Routers

- Multiple Chip Vendors
- Consistent Protocols
- Consistent App interface

Any User expectation for managing servers should be met for network devices





2015-2016 Features in linux (just to make a point)

Kernel Additions

- Nftables
- EBPf
- TC hw integration
- New Bridge driver
- VXLAN driver enhancements
- VRF
- LWT (infra for MPLS)
- Link state management
- Optimize IPv4 FIB lookup, route driven congestion algorithm selection
- Switchdev support for Mellanox Switch, DSA devices
- NetCP (network coproc) driver support
- TCP fingerprinting

Userland Upgrades

- Quagga
- iproute2
- Ethtool
- msdpd
- Ildpd
- libnl
- ifupdown2



2015-2016 Features in linux (just to make a point)

Kernel Additions

- Nftables
- EBPF

- TC hw integration
- New Bridge driver

And then there is XDP: combines with eBPF to eliminate non-standard appliances

- VRF
- LWT (infra for MPLS)
- Link state management

- Optimize IPv4 FIB lookup, route driven congestion algorithm selection
- Switchdev support for Mellanox

Switch, DSA devices

- NetCP (network coproc) driver support
- TCP fingerprinting

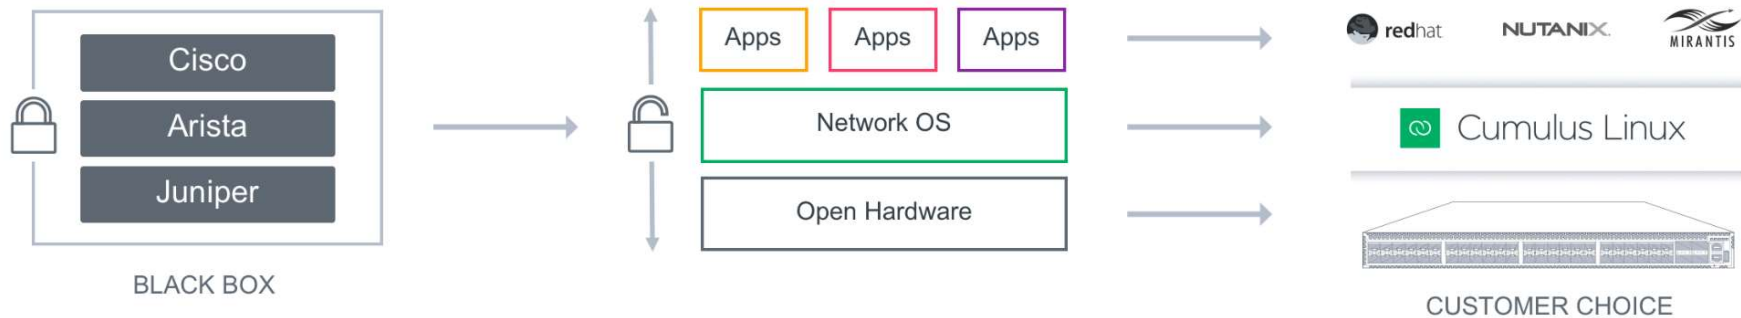
Userland Upgrades

- Quagga
- iproute2
- Ethtool
- msdpd
- Ildpd
- libnl
- ifupdown2



So .. Who is Cumulus and why does anyone care ?

Web scale networking for Enterprise Clouds





So what's a community to do

Make the package enterprise friendly



Kernel features that have major impact in the Enterprise

Ethtool

- Interface naming
- 25g, 100g standards and management

Bridging model

- Active-Active (mlog)
- IP mcast, igmp
- VXLAN scale

Route path selection

- IPv4/IPv6 uniformity

VRF

MPLS

- LWT
- MPLS datapath



But you need a greater ecosystem (driven by the kernel)

Ecosystem

- ONIE
Created the platform separation
- FRR
Routing performance, Protocols like MPLS
Capacity management => Switchdev

User Interface

- Ifupdown2
Cannot handle scale without it
- NCLU

Protocols

- EVPN
- MPLS

Highly scalable interface manager

- Supports templates
- Handle dependencies
- Validates configurations
- Extensible

Now coming to a Debian distribution near you

Search in specific suite: [\[precise\]](#) [\[precise-updates\]](#) [\[precise-backports\]](#) [\[trusty\]](#) [\[trusty-updates\]](#) [\[trusty-backports\]](#) [\[xenial\]](#) [\[xenial-updates\]](#) [\[xenial-backports\]](#) [\[yakkety\]](#) [\[yakkety-updates\]](#) [\[yakkety-backports\]](#) [\[zesty\]](#)

Limit search to a specific architecture: [\[i386\]](#) [\[amd64\]](#) [\[powerpc\]](#) [\[arm64\]](#) [\[armhf\]](#) [\[ppc64el\]](#) [\[s390x\]](#)

You have searched for packages that names contain *ifupdown2* in all suites, all sections, and all architectures. Found 1 matching packages.

Exact hits

Package ifupdown2

- [xenial \(16.04 LTS\)](#) (admin): Network Interface Management tool similar to ifupdown [\[universe\]](#)
1.0~git20151029-1: all
- [yakkety \(16.10\)](#) (admin): Network Interface Management tool similar to ifupdown [\[universe\]](#)
1.0~git20160525-1: all
- [zesty](#) (admin): Network Interface Management tool similar to ifupdown [\[universe\]](#)
1.0~git20170114-1: all

This page is also available in the following languages:

[Български \(Belgarski\)](#) [Deutsch](#) [suomi](#) [français](#) [magyar](#) [日本語 \(Nihongo\)](#) [Nederlands](#) [polski](#) [Русский \(Russkij\)](#) [slovensky](#) [svenska](#) [Türkçe](#) [українська \(ukrajins'ka\)](#) [中文 \(Zhongwen, 简\)](#) [中文 \(Zhongwen, 繁\)](#)

Content Copyright © 2017 Canonical Ltd.; See [license terms](#). Ubuntu is a [trademark](#) of Canonical Ltd. [Learn more about this site.](#)

[Report a bug on this site.](#)



What is NCLU?

A Network Command Line Utility for Cumulus Linux

- Simple to configure
- Guardrails to prevent typos & mistakes
All linux tools are preserved
- Context sensitive prompt, embedded help, and examples
No Documentation needed !!
- Rollback to a previous configuration
- Configure within bash shell
Bash access programmability



IP Routing Suite for Linux and Unix systems

for VMs, containers, switching, routing, and Internet peering

BGP4 ISIS LDP OSPF RIP

Quagga
offshoot

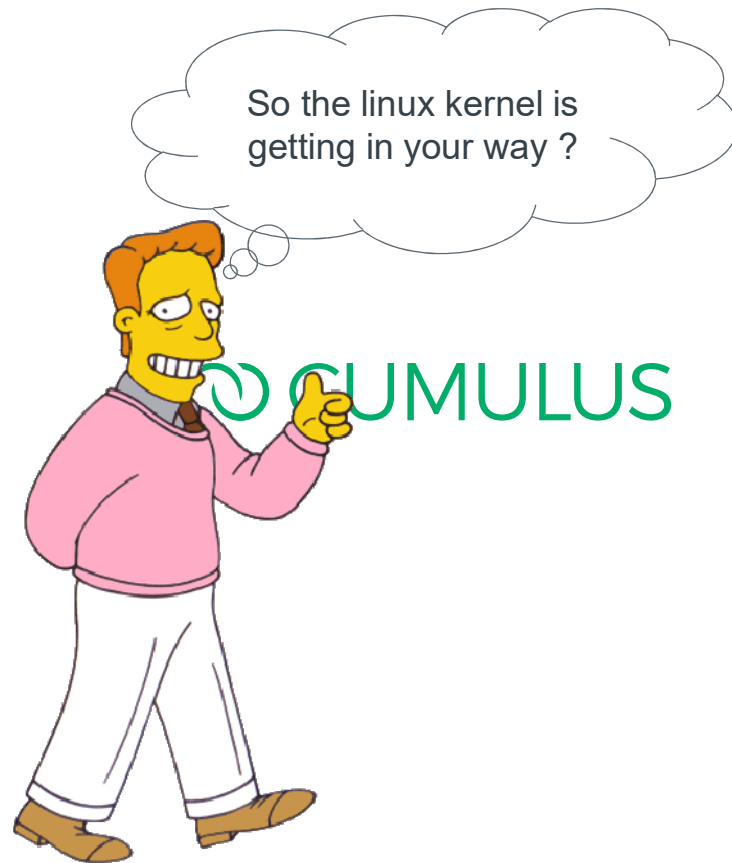
100's production
deployments

Fast, Open
Development

Community
Driven



<http://www.frrouting.org>



Finally a few words about why
Linux is the right model



Where in the stack is the right place to be innovative

Application interface: Portable and maintainable apps can use this layer

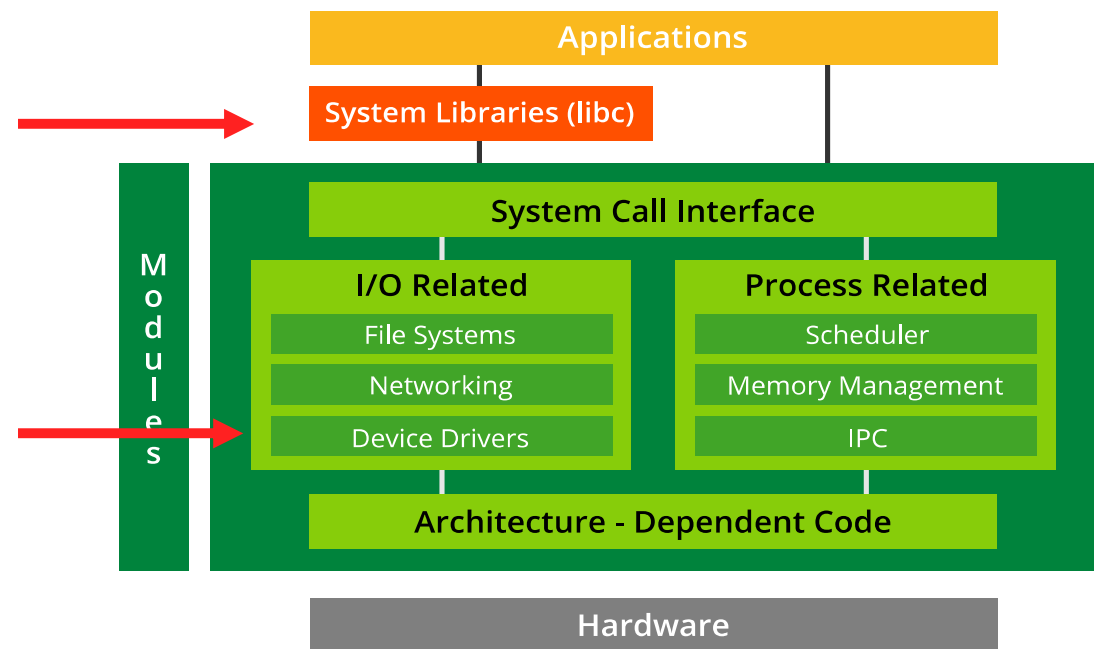
Innovation here leads to

- Non-portable apps/infra
- Highlighting narrow use cases
- Hair pulling

Low level device interface; Empowers system software developers, but only useful in the context of the full system around it

Innovation here

- Is good
- Leads to increasing improved experiences





Challenges to the model

Networking is just an App, the full stack gets in the way

- Sonic
- Openswitch
- Fboss
- Other networking stacks that use linux as a boot loader

If you use Linux, then you do not have to make the choice

- Your Hosts already run it, so tools will likely integrate better
- Switchdev, Ehtool etc formulate the layer to integrate at
- And you get the largest software community as built in support



Challenges to the model

Networking is just an App

- Sonic
- Openswitch
- Fboss

Bottom Line : Use the Linux Kernel as the model and become part of the solution



Thank you!

Visit us at cumulusnetworks.com or follow us [@cumulusnetworks](https://twitter.com/cumulusnetworks)

© 2017 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.



Cumulus Networks' customers

28%
of Fortune
50 Companies

25%
of Mega-Scale
Public Clouds

>2M
Ports Powered by
Cumulus Technology

600+
Customers Across
Many Industry Verticals

