

# OPEN

Compute Summit

January 28–29, 2014 San Jose





# Open Network Linux

## A Common Linux Platform for OCP Switches

Rob Sherwood  
Big Switch Networks  
CTO

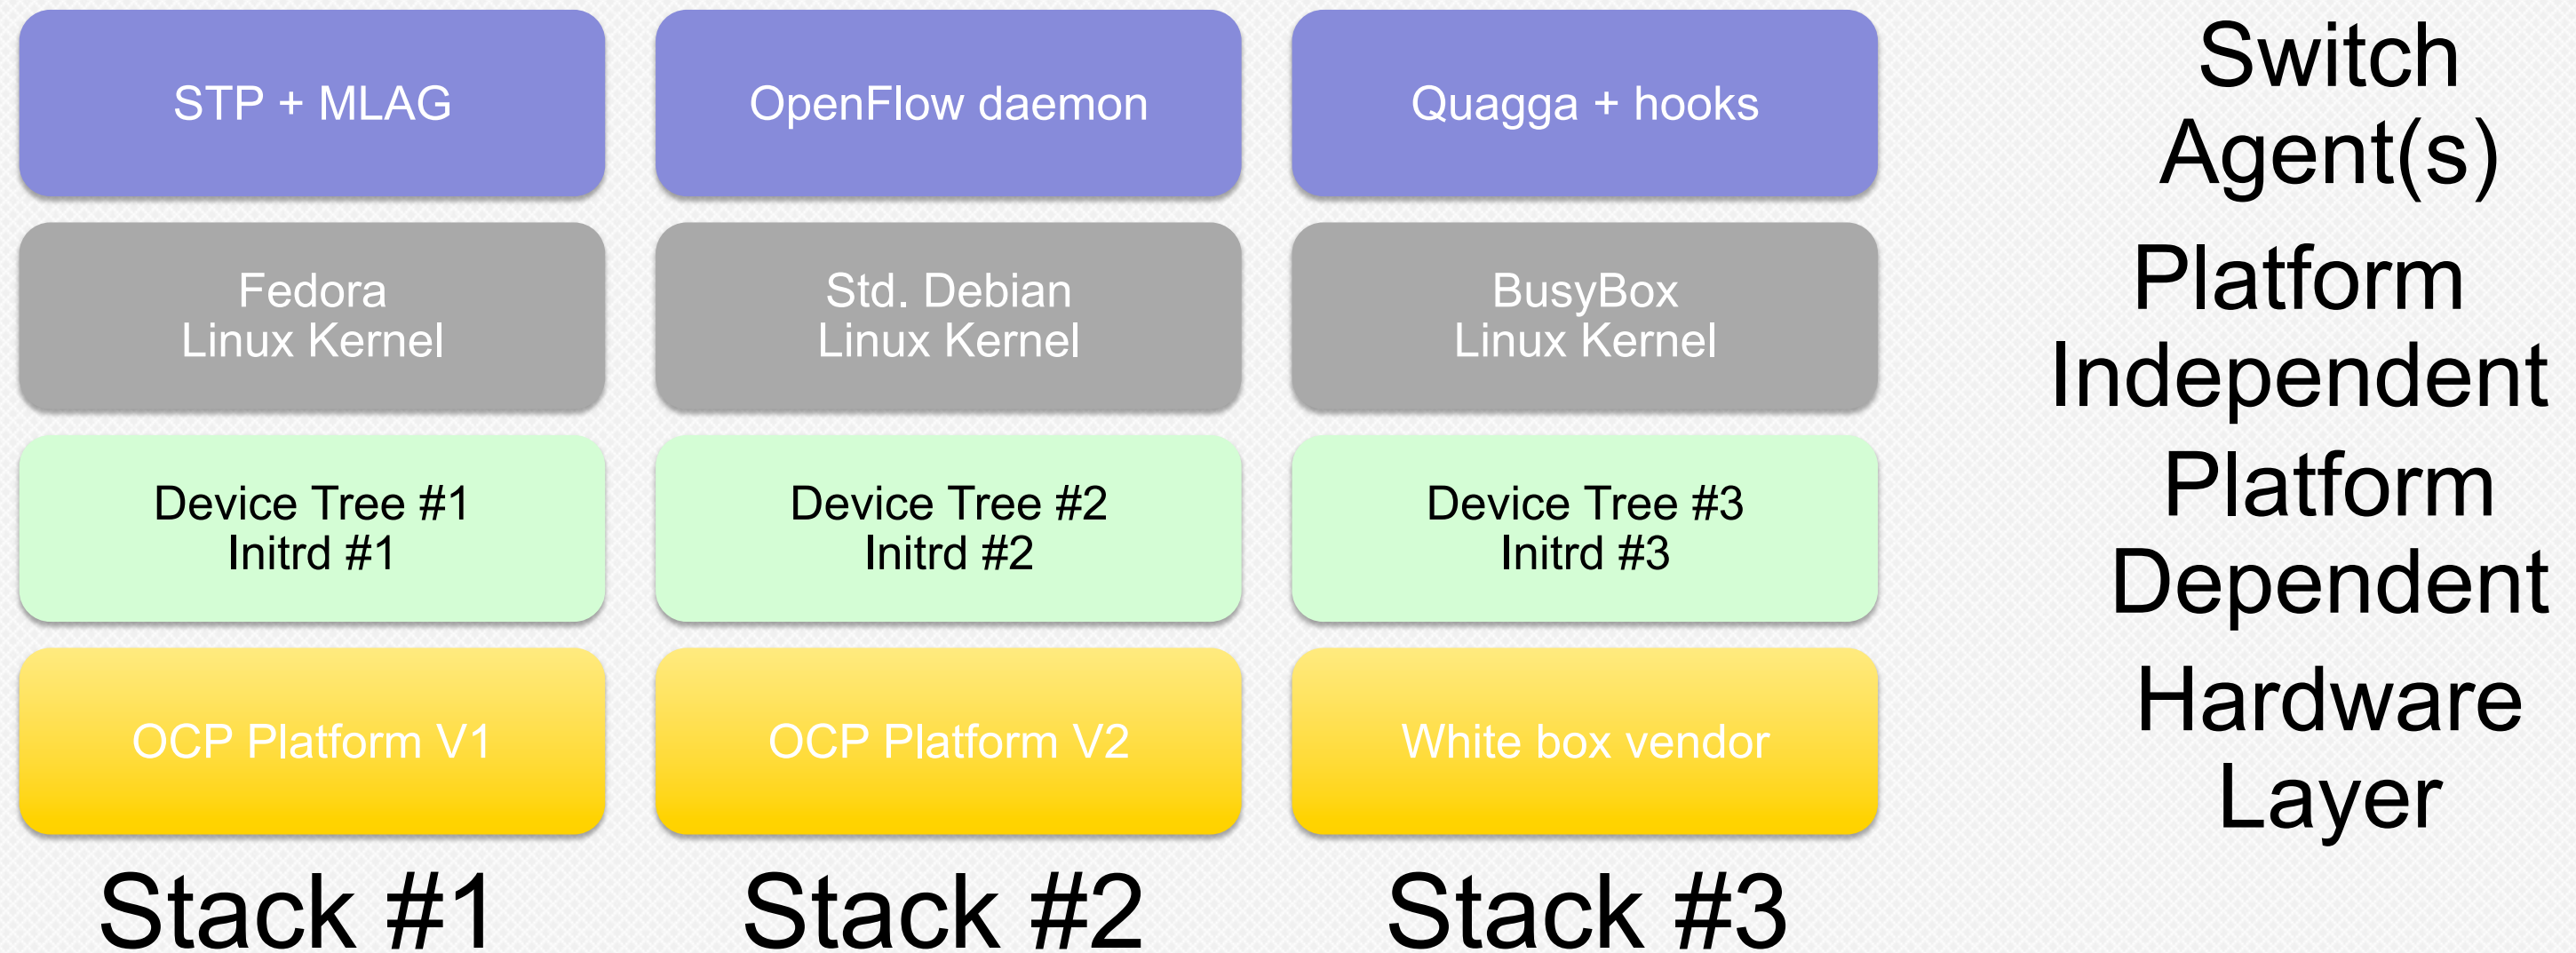


# Outline

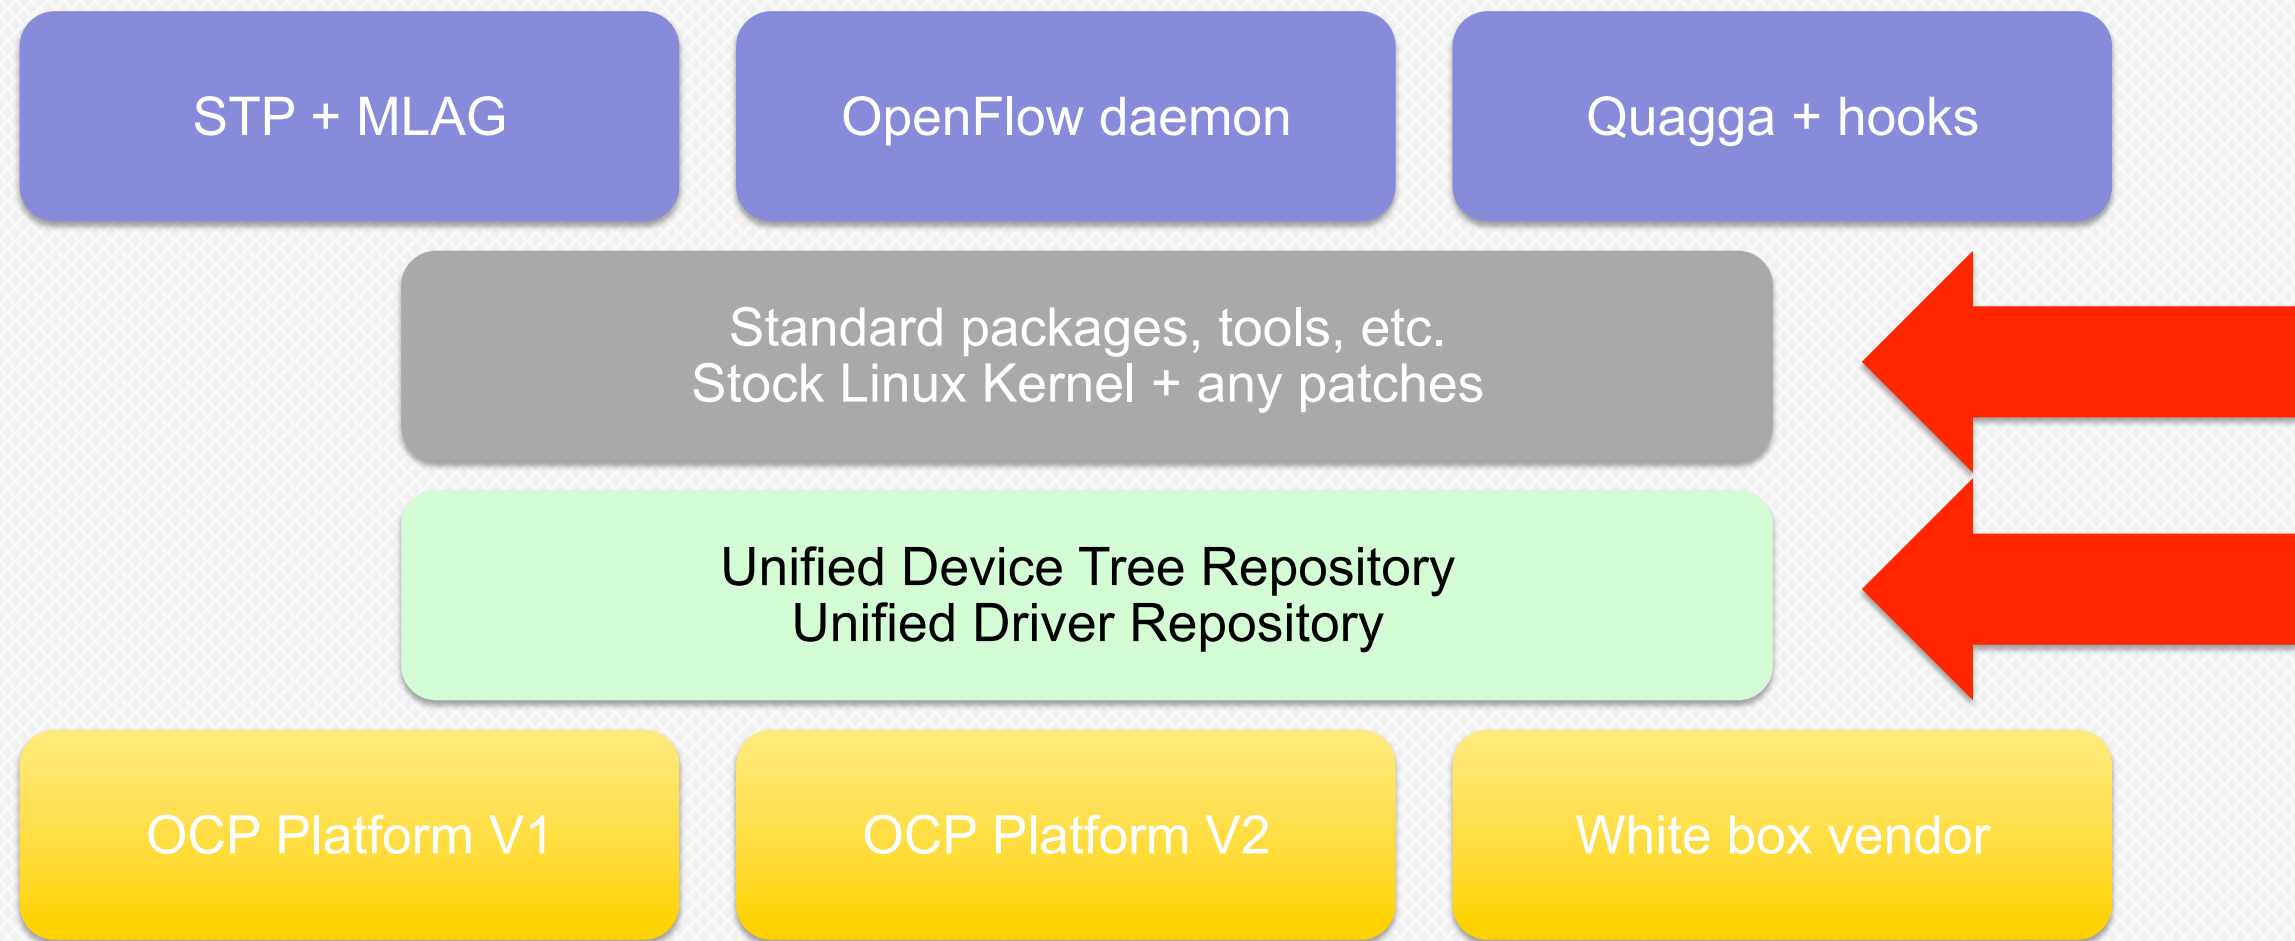
- Proposed in November OCP workshop
  - Goal: Common community target → faster adoption
  - Status: Open sourced on github yesterday: 1/27 (!!)
- Technical details
  - Multi-platform support: x86, PPC, and x86 VM
  - Full “Server-like” experience on network hardware
  - Network booting and image management
- Demo and “Hello World” L3 forwarding app



# November Proposal: Tower of Babel is Bad



# November Proposal: Common Linux Platform



Keep  
differentiation in  
switch agents

Come together  
around the  
common bits

Maximize  
hardware  
abstraction



# Open Network Linux: *Goals*

- **Accelerate adoption of OCP switch hardware**
  - Users: download image, install via ONIE
  - Vendors: common Linux platform for new drivers, testing
- **Create an open community**
  - Target: Linux portable to all networking devices
- **License: Eclipse Public License and GPL for Kernel**
- **“What’s in it for me?”**
  - Engineering efficiencies
  - Better development and deployment experience



# Open Network Linux: *Status*

- Lots of support from community – thanks!
- Github went live yesterday: 1/27/2014
  - Main repository: [github.com/opennetworklinux/ONL](https://github.com/opennetworklinux/ONL)
- Builds ONIE-compatible images for:
  - Generic x86 platforms: Interface Masters not *yet* tested
  - Many PPC platforms: Quanta LY2, LB9, **LB8D**, **Accton 5652**
  - x86 VM build: for testing – qcow2 (vmdk via convert)
- Stability level: “works for us”
  - Feedback welcome



# Outline

- Proposed in November OCP workshop
- Goal: Common community target → faster adoption
- Status: Open sourced on github yesterday: 1/27 (!!)
- **Technical details**
  - Multi-platform support: x86, PPC, and x86 VM
  - Full “Server-like” experience on network hardware
  - Network booting and image management
- **Demo and “Hello World” L3 forwarding app**





# Technical Overview

- **Code builds two main artifacts:**
  - ONL installer/loader: like grub, but multi-platform with netboot
  - ONL SWI file: zip'd **SW**itch *I*mage with root fs, kernel, initrd
- **Code is divided into multiple sub-modules**
  - *ONL*: main repository – auto pulls in other repos
  - *linux-3.9.6*: extracted kernel code
  - *loader*: scripts and code for boot loader process
  - *infra* and *common*: libraries, shared routines, faultd



# Deployment Overview

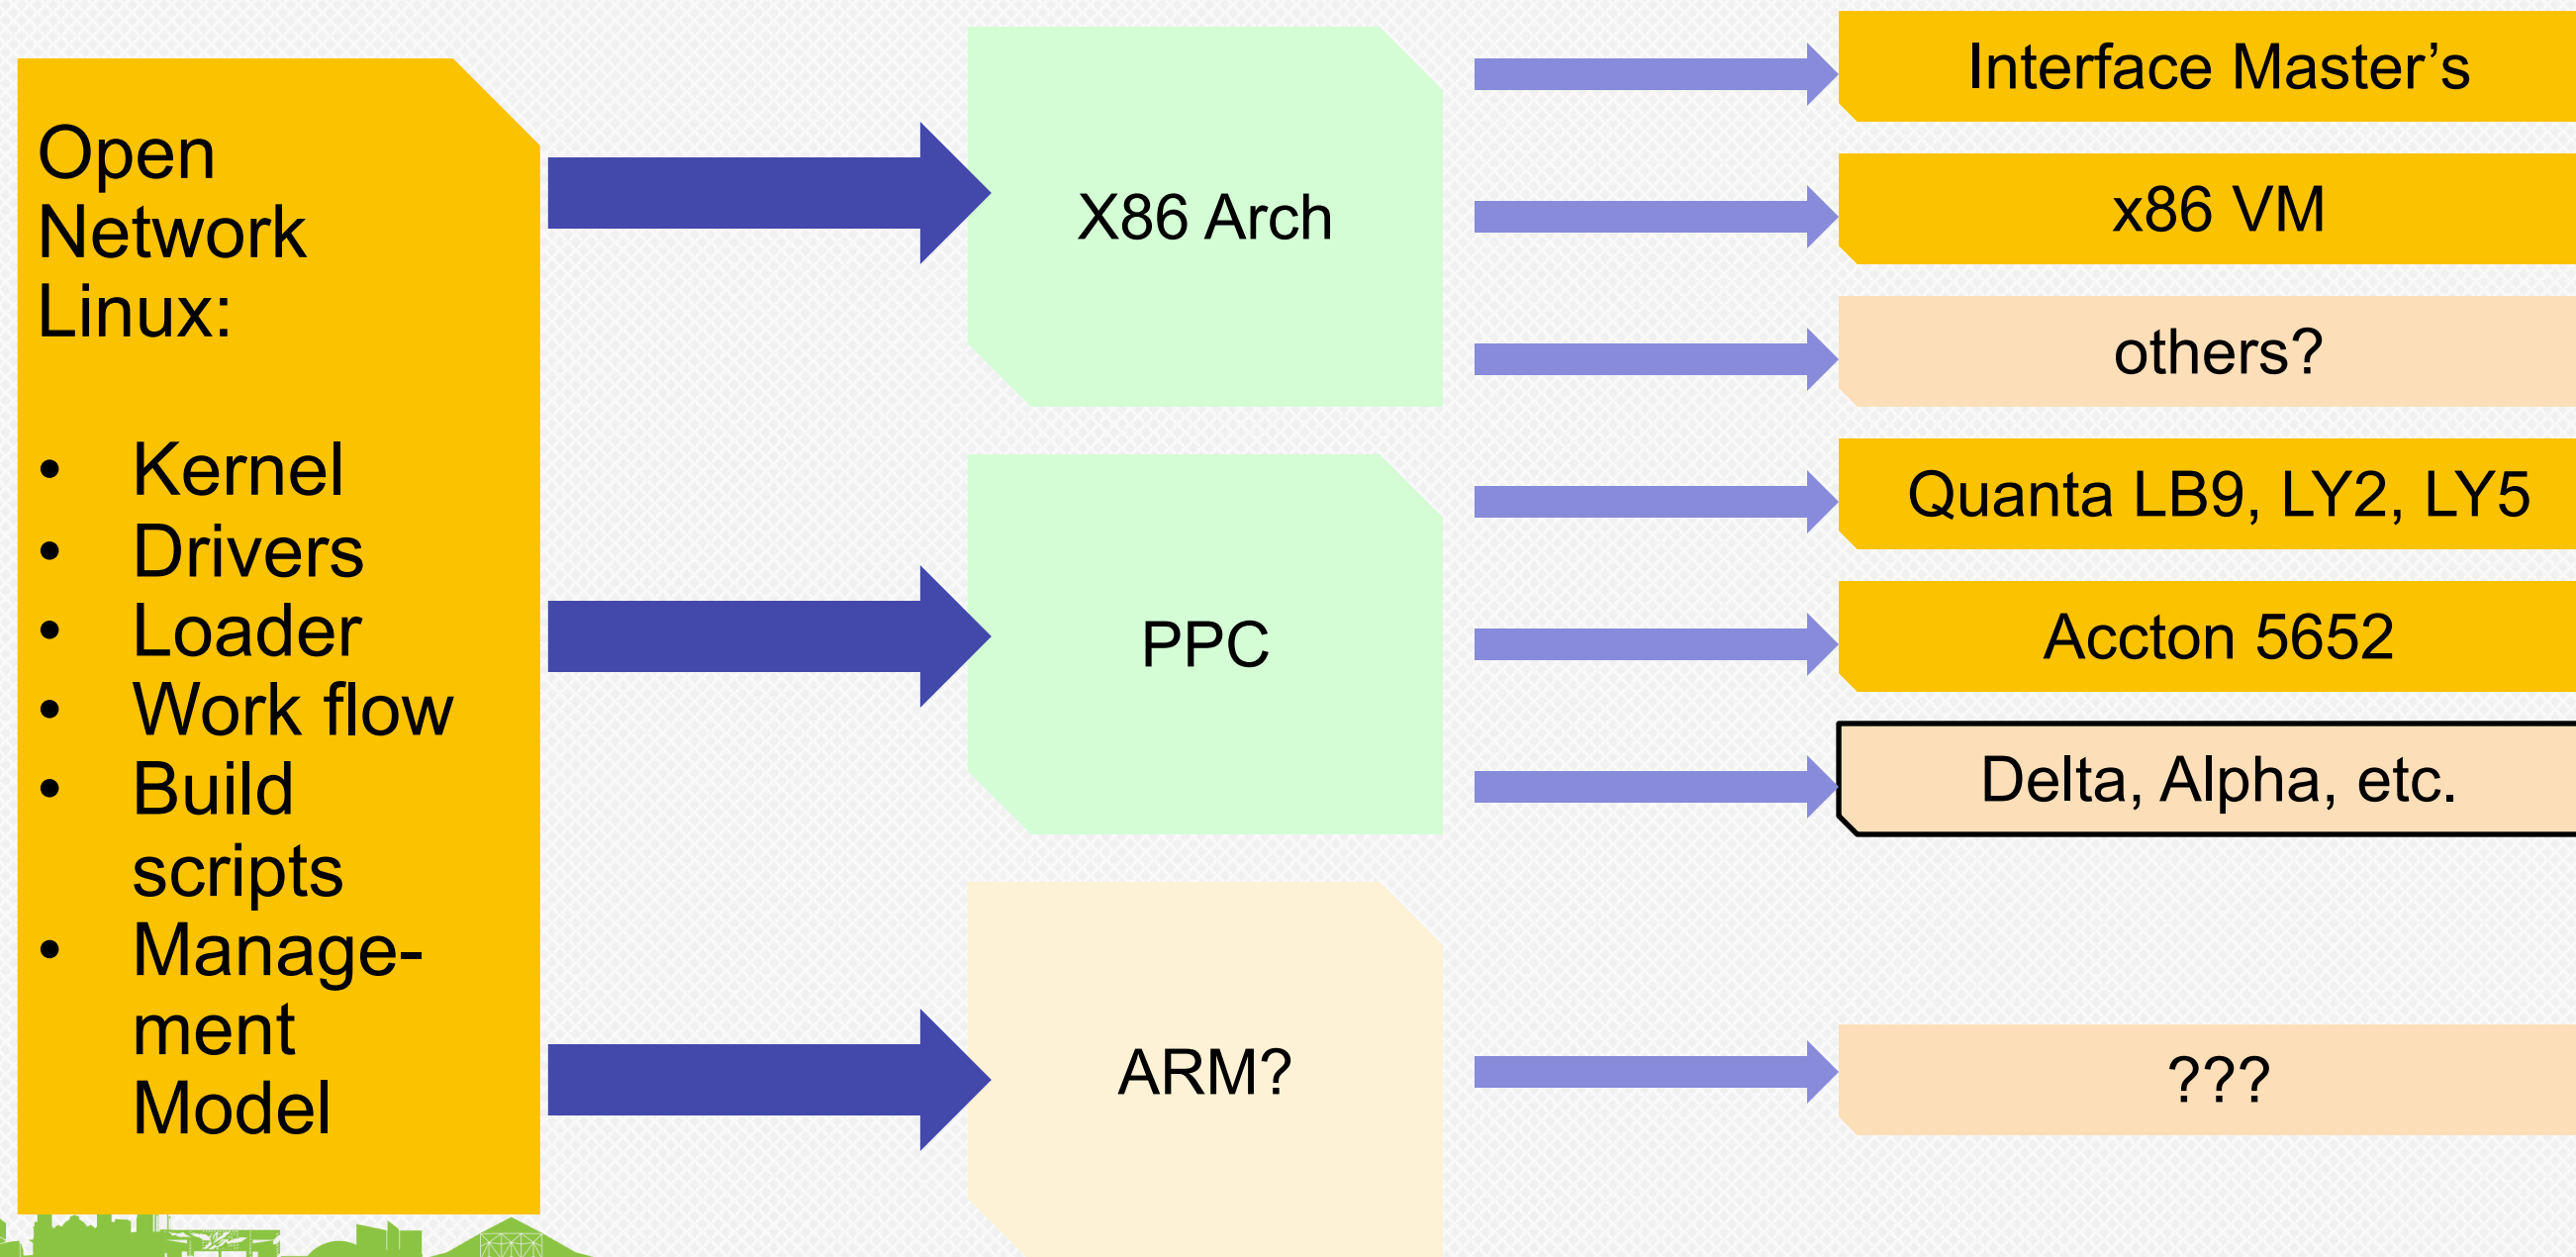
- Full documentation in README in ONL repository

1. Get code from [github.com/opennetworklinux/ONL](https://github.com/opennetworklinux/ONL)
2. (only for ppc) Build a cross-compilation workspace
3. Build ONL installer/loader image
4. Put ONL installer/loader image on ONIE server
5. Boot switch and install ONL via ONIE
6. Build one or more ONL SWI's
7. Netboot from scp/nfs/http/ftp/etc. to install ONL SWI



# ONL is Multi-Platform

- Support many boxes from the same code-base



# Tricks to Use Switches Like Servers

- **Switches have flash, not hard drives**
  - Problem 1: Maximum flash cycle time limit disk writes
  - Problem 2: Flash and ram more limited than typical servers
  - Fix: Use overlayfs to overlay copy-on-write ram disk over flash
- **ONL uses full-featured binaries**
  - For size, most switch OS's use stripped binaries, e.g., busybox
  - Bigger binaries uses additional space, but ok with overlayfs
  - Install/use proper Debian binaries using apt-get
  - Useful for development or operations, e.g., gcc or Chef/Puppet



# ONL Supports Net Booting Natively

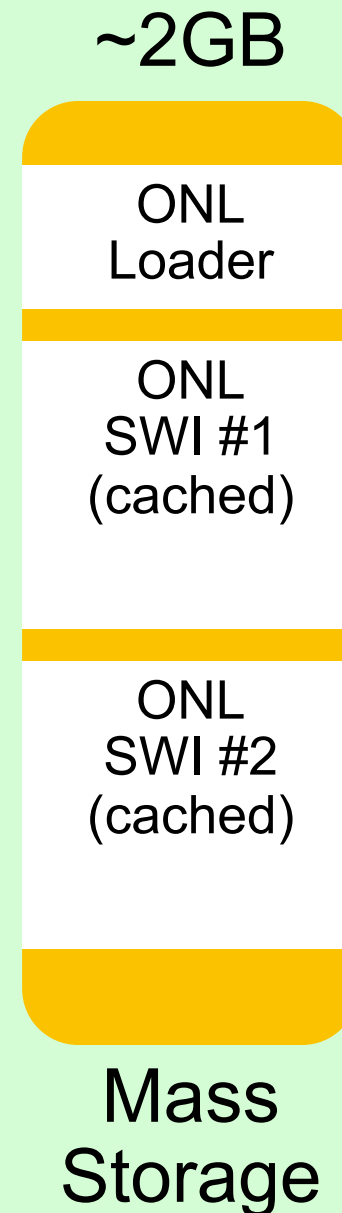
- **Boot SWIs over the network once ONL Loader installed**
  - SWIs are cached locally for performance, resilience
  - Simplifies operational management, upgrades
  - Supports http, ftp, tftp, nfs, ssh/scp, or **ZTN**
- **Zero-touch Networking (ZTN): auto-discover SWIs**
  - Like PXE for your switches
  - Query all http servers in local subnet, use SWI of first hit
  - Just like ONIE
- **Netboot makes managing your network *much* easier**



# Install Using ONIE then Boot ONL

## Boot Logic:

1. uBoot POSTs
2. \$nos\_boot\_cmd is read from ENVs
3. run \$nos\_boot\_cmd
  - If \$nos\_boot\_cmd returns, run ONIE
  - On install, ONIE sets \$nos\_boot\_cmd to load ONL loader
4. Loader downloads specified SWI URL if not cached
5. Loader mounts rootfs as ramdisk with overlays
6. ONL loader kexec's SWI kernel



# Outline

- Proposed in November OCP workshop
- Goal: Common community target → faster adoption
- Status: Open sourced on github yesterday: 1/27 (!!)
- Technical details
  - Multi-platform support: x86, PPC, and x86 VM
  - Full “Server-like” experience on network hardware
  - Network booting and image management
- Demo and “Hello World” L3 forwarding app



# Demo and Forwarding Agents

- **Demo:**

- Netbooting
- Quick file system walk through

- **Available Forward agents**

- Download BSN's Switch Light OpenFlow daemon binary
- "Hello World" L3 agent
  - Monitor Linux software route via rtnetlink
  - Copy routes/neighbor/interfaces into ASIC with binary driver
  - Enables box for Quagga, Xorp, etc.
  - *Work in progress*





# Conclusion

- **Open Network Linux is available now**
  - Goal: support all bare metal switches: OCP and non-OCP
  - [github.com/opennetworklinux/ONL](https://github.com/opennetworklinux/ONL)
- **Technical benefits:**
  - Multi-platform, switch like server, netbooting
- **Contributions Encouraged!**
  - New platforms
  - Additional features/drivers
  - Better documentation





This is an interstitial slide

