OPEN Compute Summit

Open Network Linux

A Network Operating System (NOS) for OCP

Rob Sherwood
Big Switch Networks
CTO

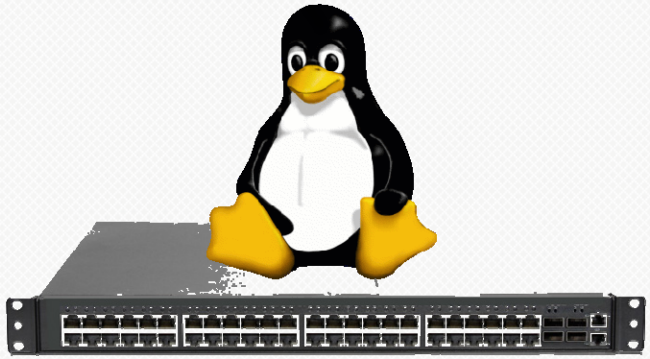

Engineering Workshop

Outline: Open Network Linux (ONL)

- What is ONL?
- Contributors and Community
- What does ONL provide?
- ONIE Compatible NOS Installer
  - Second stage more featureful loader
  - “Server-like” Management
  - Platform Abstraction Layer – ONLP
  - Supported Hardware

What is Open Network Linux?

- A collection of software packages, utilities, drivers, and abstractions to run on OCP Switch hardware
  - i.e., a “NOS” that ONIE would install
- Why not use an existing Linux distribution?
  - Does build on existing distribution – Debian Wheezy
  - Need to create ONIE installers for many platforms
  - Need to manage switch-specific hardware (e.g.,SFPs)
  - Switches are very similar to servers, but not quite

ONL Example Use-Cases

1. DIY packet forwarding platform, e.g., for students
   - Platform driver code too complicated for most hobbyists
2. OCP Certification Reference testing platform
   - Common core in open source for verification, reference
3. Building Block for Commercial Software
   - Pull ONL in to a larger, commercially supported solution

Anatomy of `make onl-$arch`

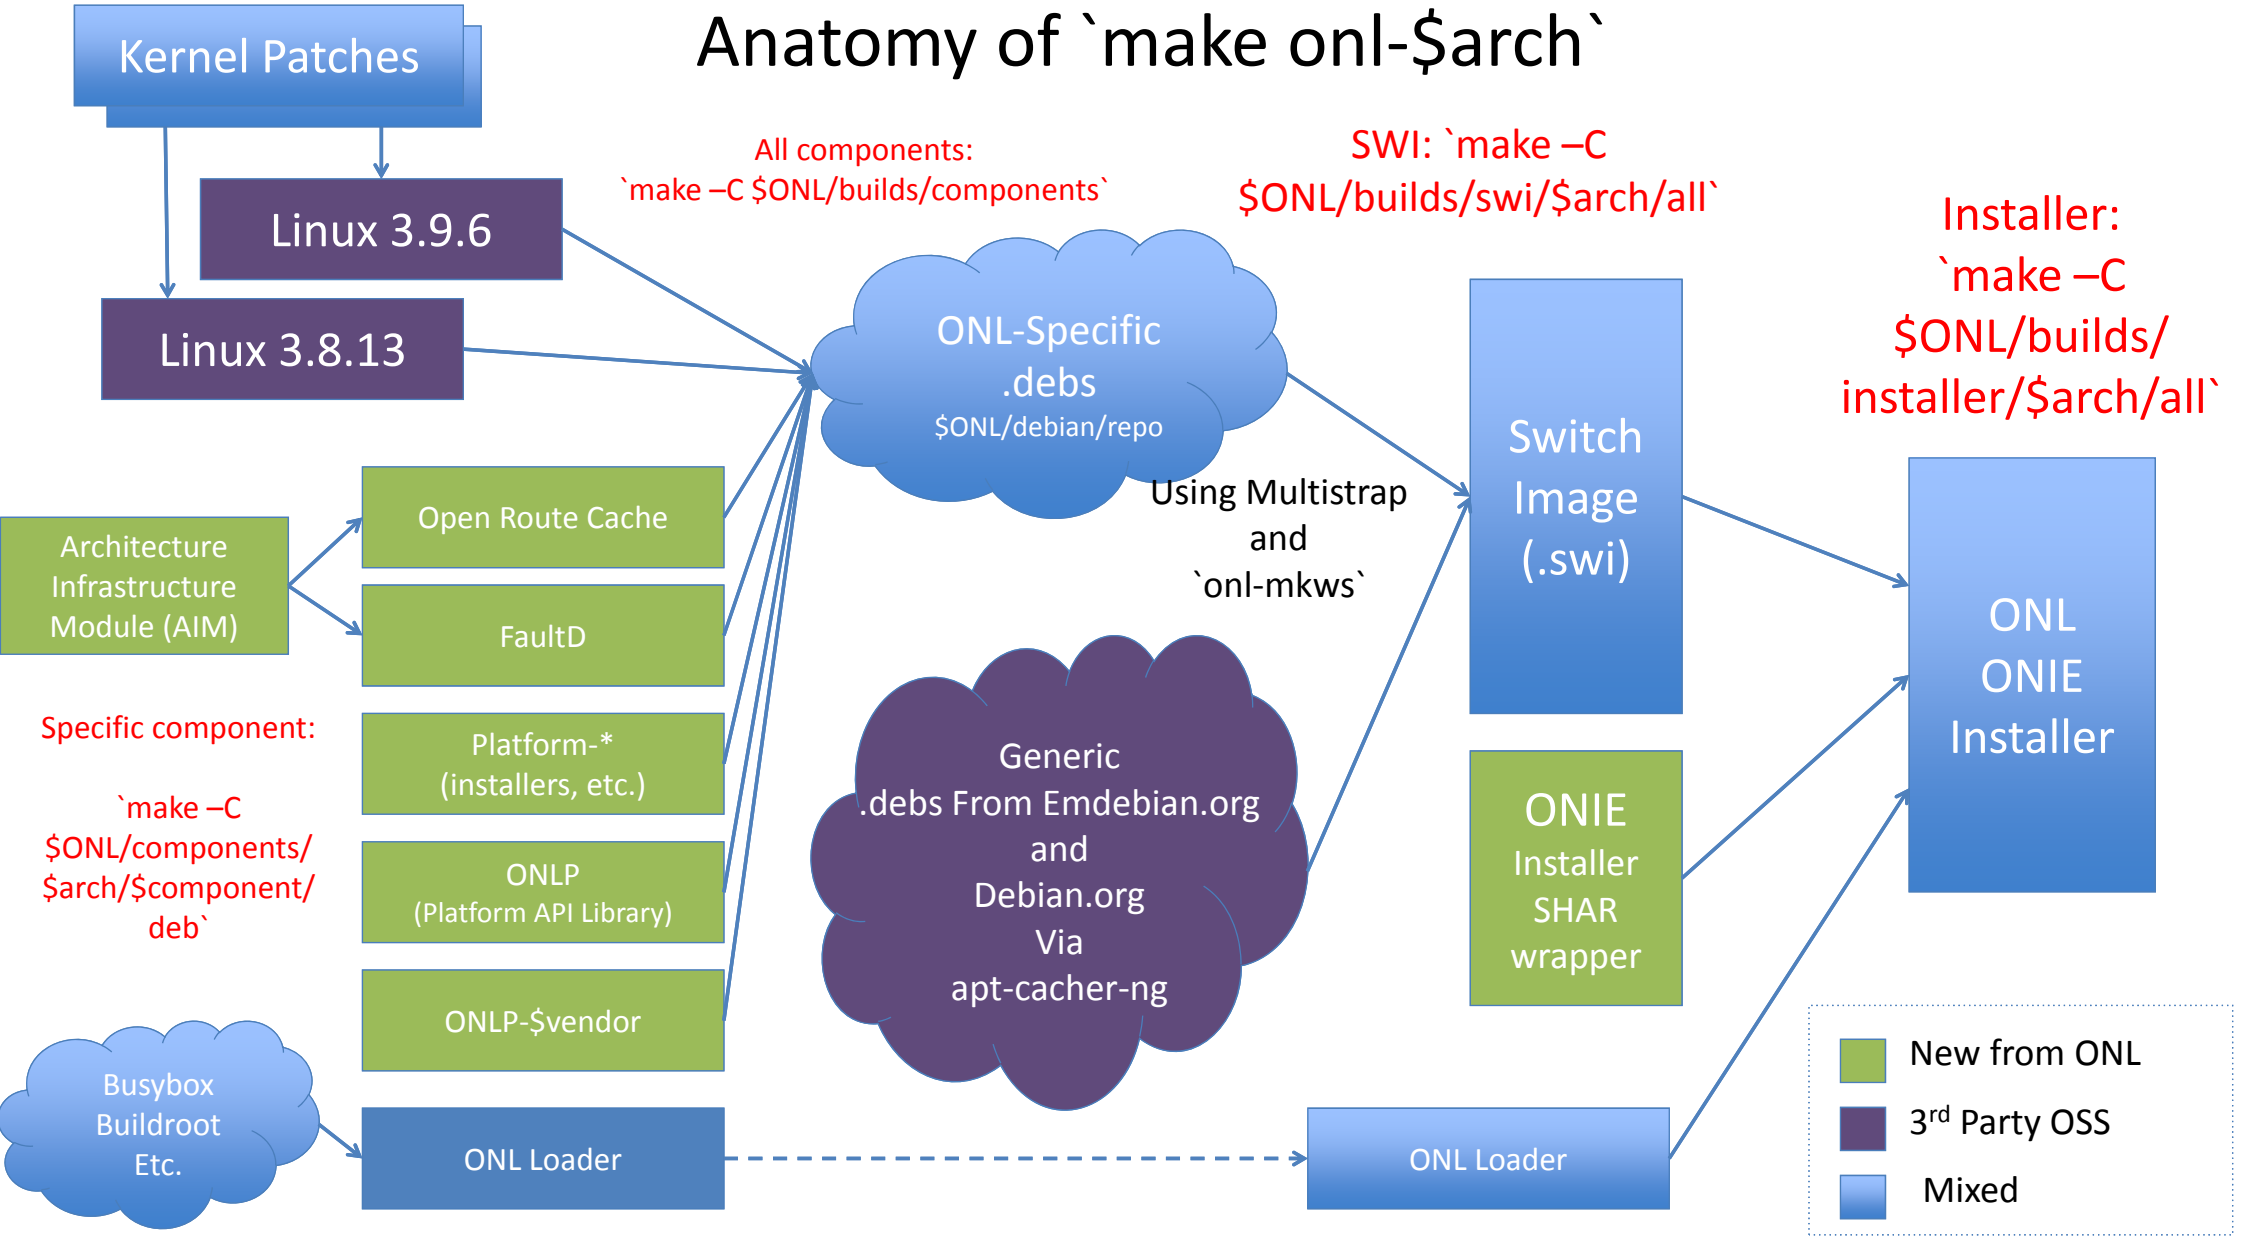

ONL Relative to ONIE

"An Installer Needs a NOS"

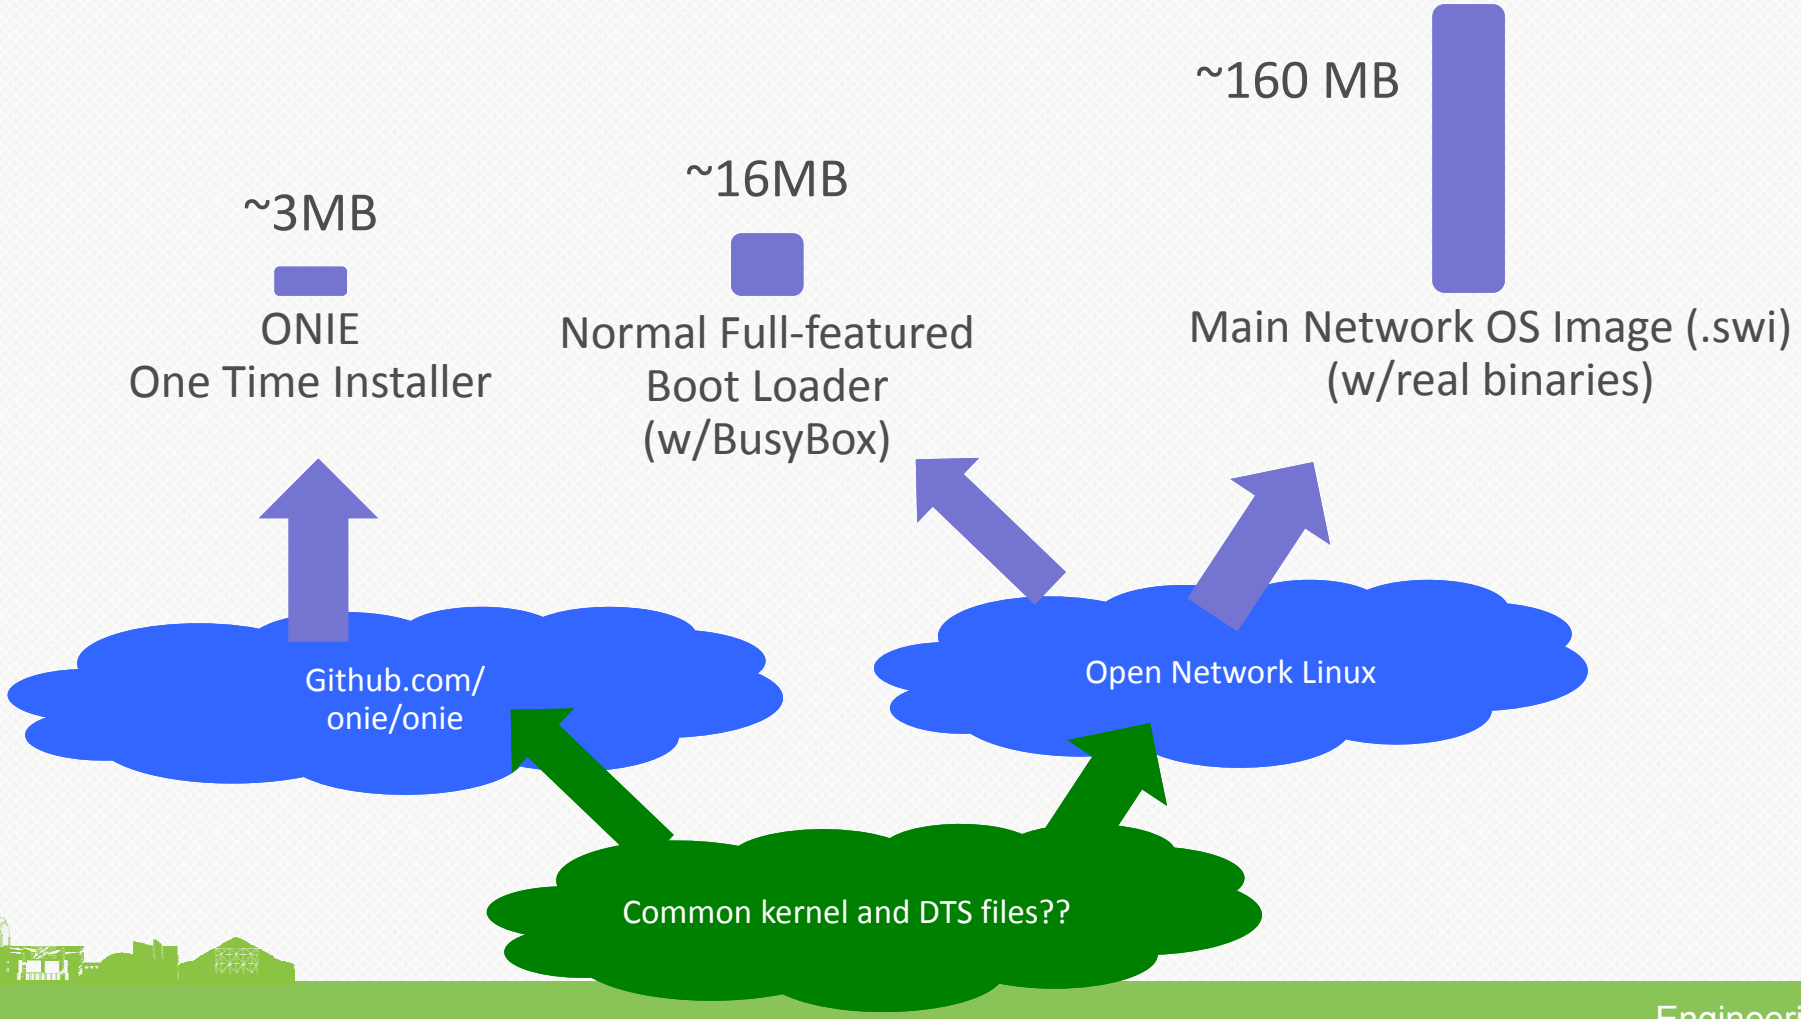

Why Use ONL?

- Help ecosystem focus on innovation
  - Many annoying software details to run an OCP switch
  - Building platform drivers not high value asset; should be common
- Enables a reference NOS implementation
  - Hardware without software is not useful
  - Package up details and best practices into one place
- Help bootstrap the Open ecosystem and OCP adoption
  - Allows commercial companies and DYI-folks to build OCP-based products faster

Outline: Open Network Linux (ONL)

- What is ONL?
- Contributors and Community
- What does ONL provide?
- ONIE Compatible NOS Installer
  - Second stage more featureful loader
  - “Server-like” Management
  - Platform Abstraction Layer – ONLP
  - Supported Hardware

Background: Generic NOS Architecture

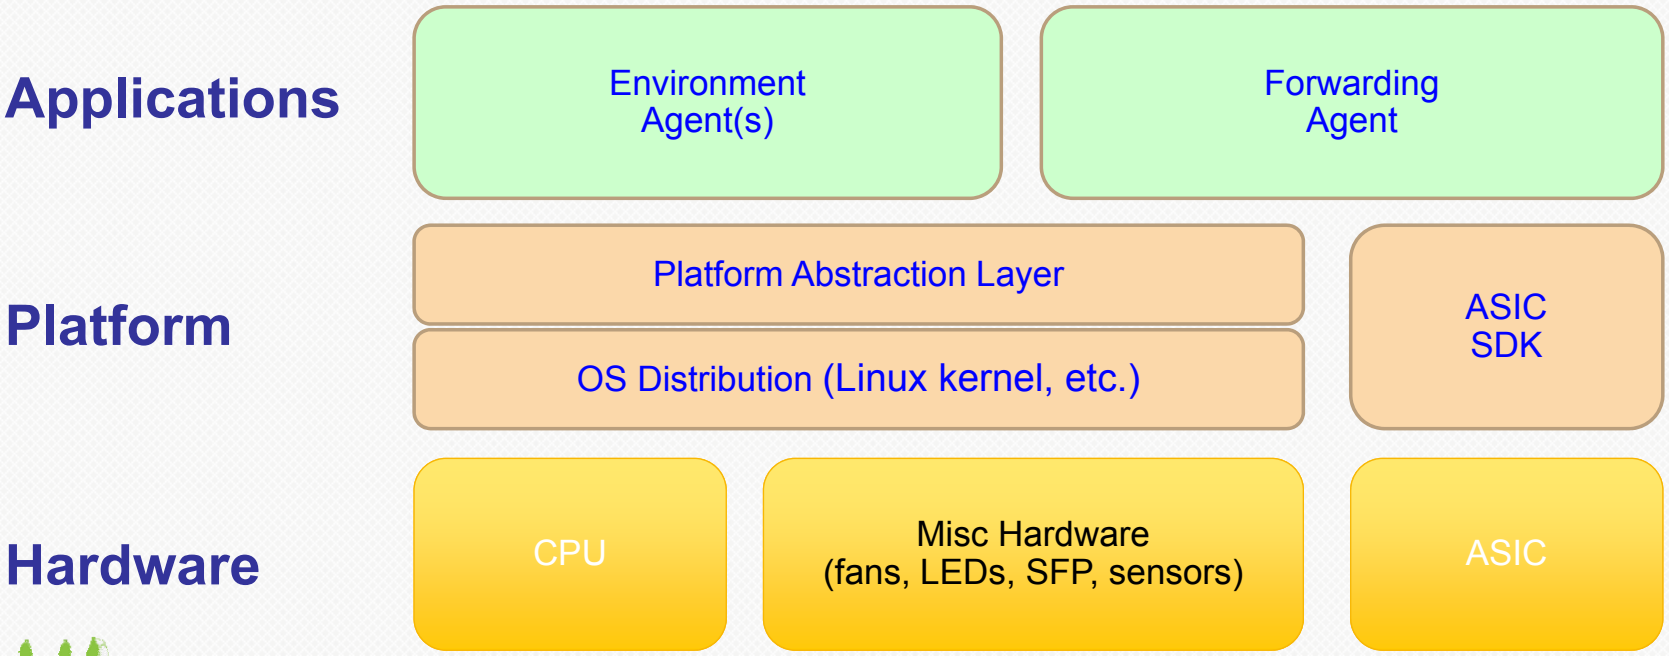

Shipping now: ONL + Open Route Cache

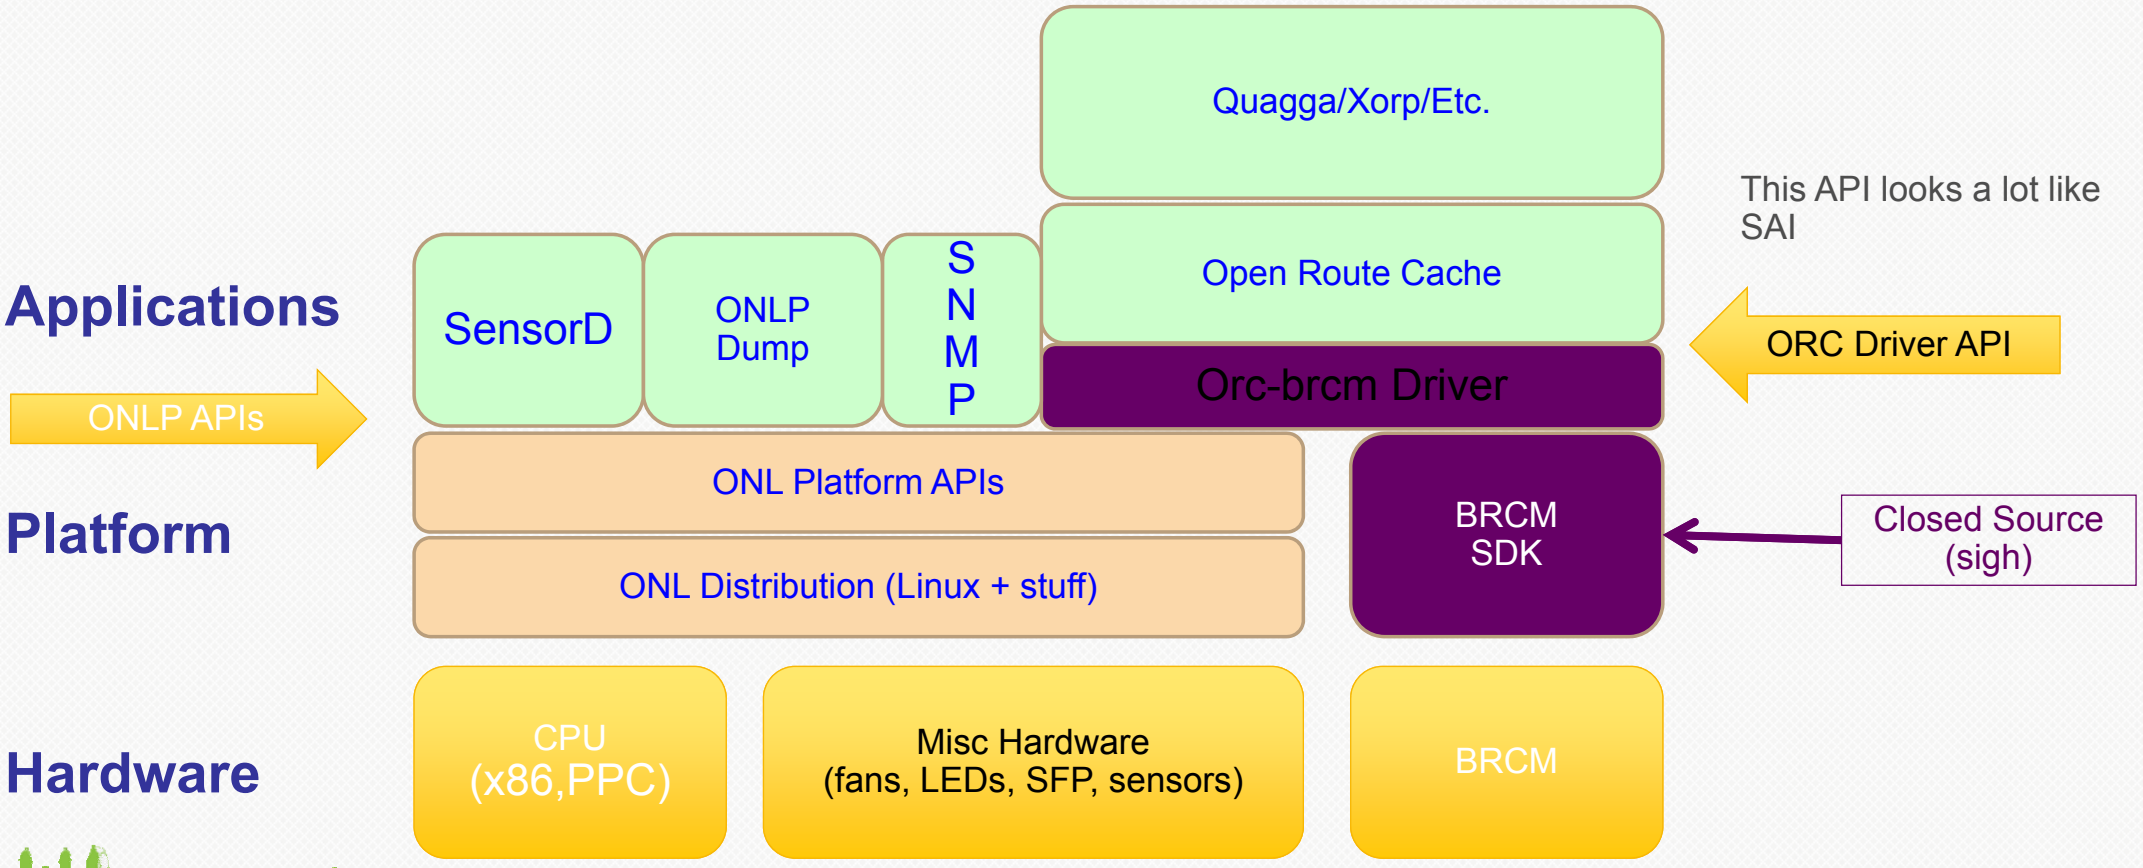

Per Switch Stack

Next Target: ONL + OpenFlow Agent

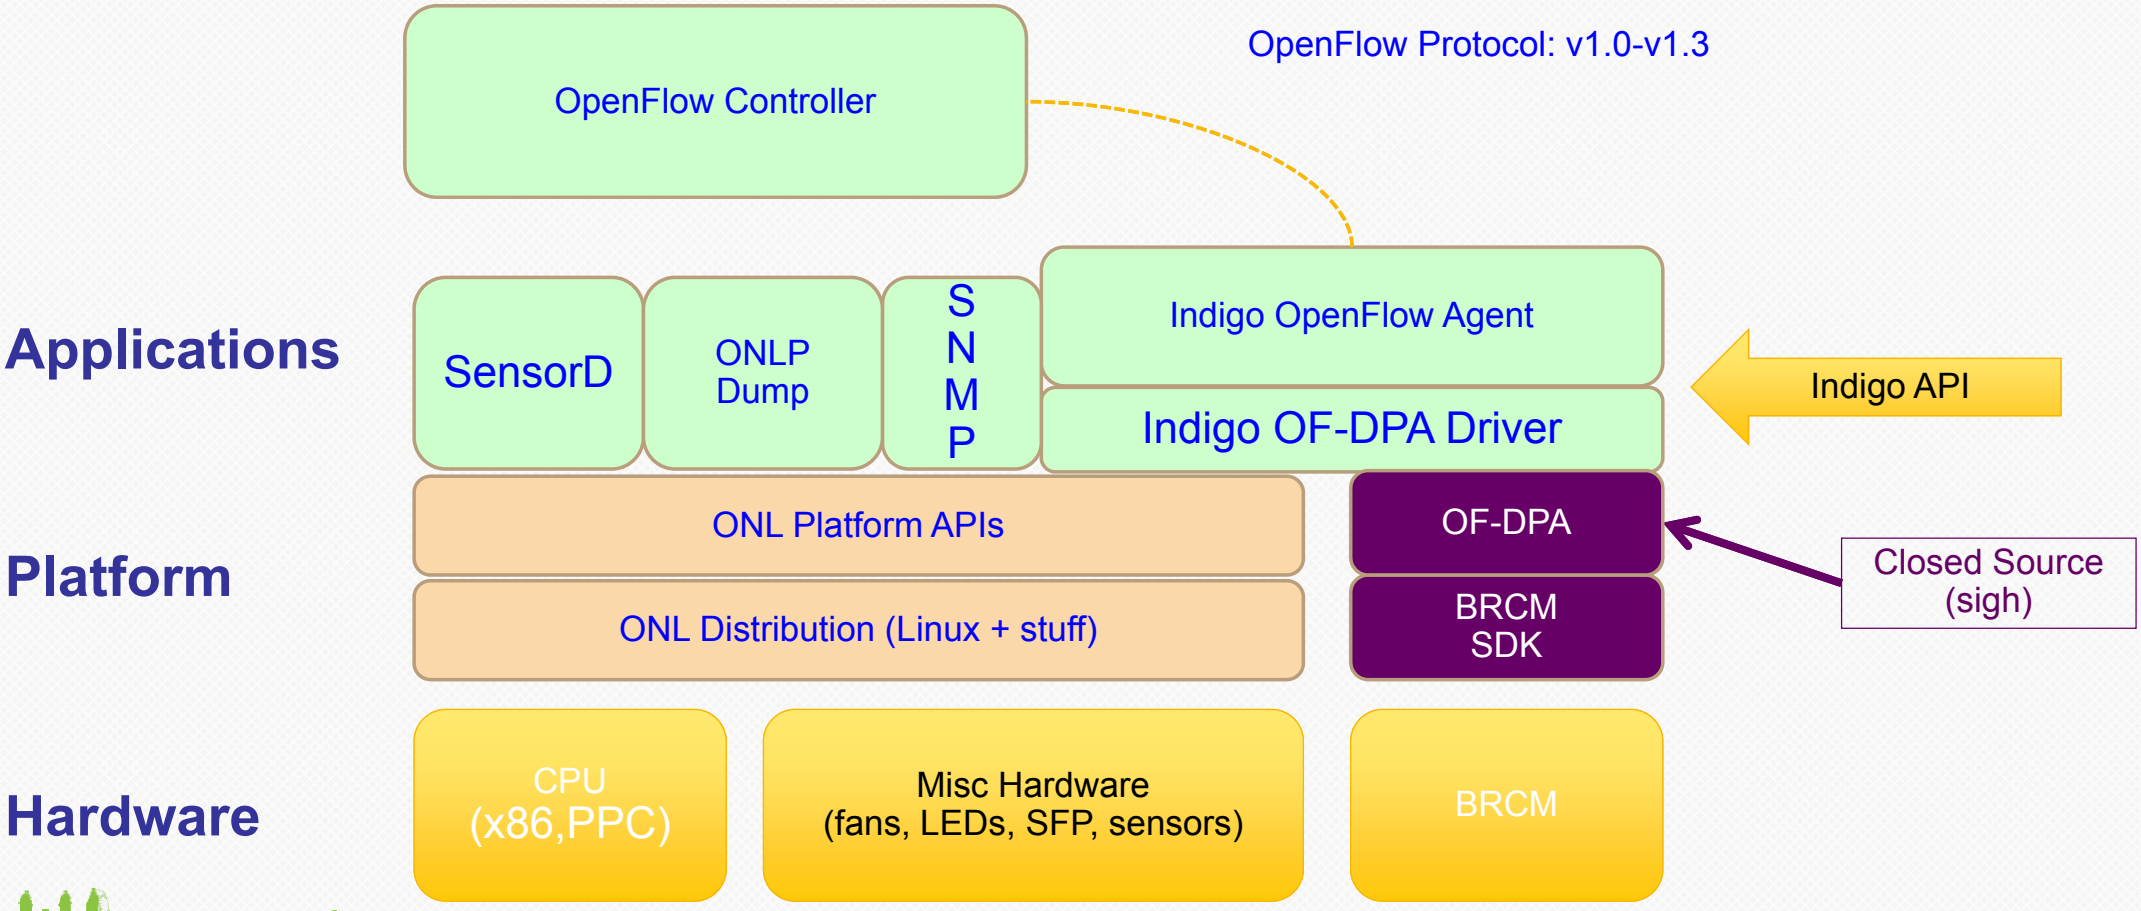

ONL Contributors/Supporters (So Far)

Adds Forwarding Agents

big switch networks

PICA8 OPEN NETWORKING

Open Network Linux

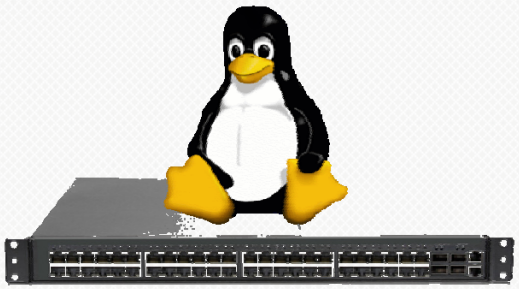

Finisar

freescale

Additional Drivers

Provides Platform Drivers

Accton

Interface Masters TECHNOLOGIES Innovative Network Solutions

OTHERS

Open Network Linux Community

- Project has been public for 10 months
  - Code – http://github.com/opennetworklinux/ONL
  - Website – http://opennetlinux.org
  - ~50K LOC (without Linux patches); 56K Makefile lines
- Community Traction
  - Five six platforms added by community; more coming
  - Three companies have commercial projects w/ONL
  - ~14 code contributors from four companies

Outline: Open Network Linux (ONL)

- What is ONL?
- Contributors and Community
- What does ONL provide?
- ONIE Compatible NOS Installer
  - Second stage more featureful loader
  - “Server-like” Management
  - Platform Abstraction Layer – ONLP
  - Supported Hardware

ONL Provides an ONIE-Compatible Installer

- ONIE is great – but needs an installer
  - Installer packages entire NOS into a single image
- Installer is the glue between Platform and NOS
  - Sets the $nos_bootcmd magic
  - e.g., nos_bootcmd= diskboot 0x10000000 0:1 ; bootm 0x10000000
  - Formats local storage, e.g., flash or local disk
  - Installs ONL loader onto local storage (see next)

ONL Provides A Second-Stage Loader

- uBoot (PPC) or Grub (x86) is the first stage loader
  - ONIE is a 2nd-stage loader; only runs at install-time
- ONL 2nd Stage Loader
  - Runs every time before NOS boots
  - Built on Linux – enables a full shell for rescue mode
  - Load NOS via ssh/scp/http/tftp/nfs/ftp/local flash
    - Net boot allows centralized NOS image management
  - (future) Initialize ASIC/front panel ports – inline boot

ONL Provides “Server-like” Management

- 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 has a Platform Abstraction Layer

- BIOS does not provide inventory/drivers for all devices
  - SFP has a life cycle, like USB; F2B vs. B2F power supplies
  - Hot Pluggable devices: PSU, Fan trays, AC/DC power
- Not all platforms have BIOSs
  - Need for per-platform memory maps of LEDs, Fans, Temp
- ONL Platform provides an Platform Abstraction Layer
  - Not tied to Linux subsystem (but could be integrated)
  - Goal is for trivial driver impls; leverage existing diag code
  - OID based; single threaded, multi-application
  - https://github.com/opennetworklinux/ONLP/tree/master/modules/onlp/module/inc/onlp

ONL – Supported Hardware

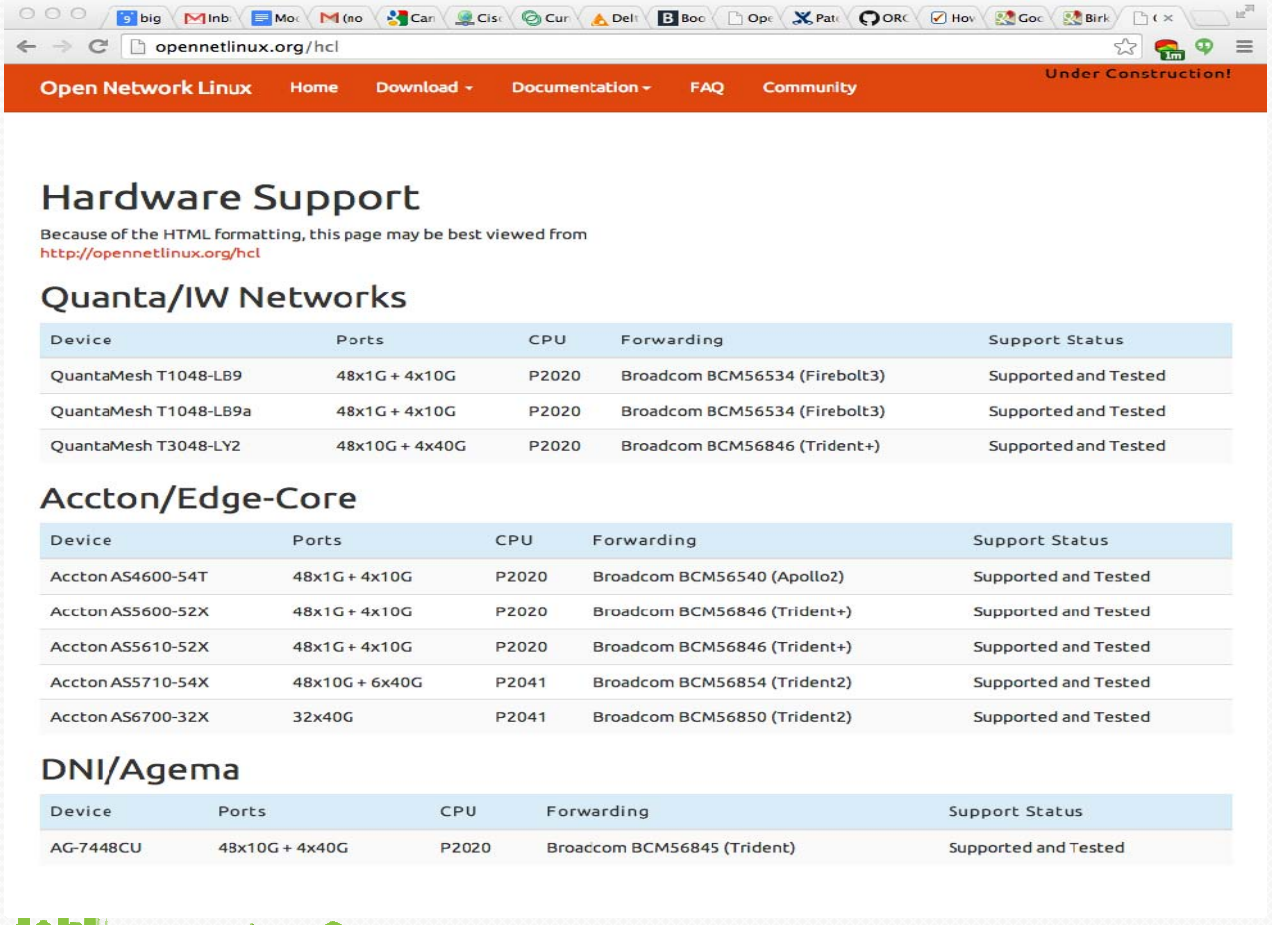

opennetlinux.org/hcl

Open Network Linux | Home | Download | Documentation | FAQ | Community

Under Construction!

Hardware Support

Because of the HTML formatting, this page may be best viewed from
http://opennetlinux.org/hcl

Quanta/IW Networks

Device	Ports	CPU	Forwarding	Support Status
QuantaMesh T1048-LB9	48x1G + 4x10G	P2020	Broadcom BCM56534 (Firebolt3)	Supported and Tested
QuantaMesh T1048-LB9a	48x1G + 4x10G	P2020	Broadcom BCM56534 (Firebolt3)	Supported and Tested
QuantaMesh T3048-LY2	48x10G + 4x40G	P2020	Broadcom BCM56846 (Trident+)	Supported and Tested

Accton/Edge-Core

Device	Ports	CPU	Forwarding	Support Status
Accton AS4600-54T	48x1G + 4x10G	P2020	Broadcom BCM56540 (Apollo2)	Supported and Tested
Accton AS5600-52X	48x1G + 4x10G	P2020	Broadcom BCM56846 (Trident+)	Supported and Tested
Accton AS5610-52X	48x1G + 4x10G	P2020	Broadcom BCM56846 (Trident+)	Supported and Tested
Accton AS5710-54X	48x10G + 6x40G	P2041	Broadcom BCM56854 (Trident2)	Supported and Tested
Accton AS6700-32X	32x40G	P2041	Broadcom BCM56850 (Trident2)	Supported and Tested

DNI/Agema

Device	Ports	CPU	Forwarding	Support Status
AG-7448CU	48x10G + 4x40G	P2020	Broadcom BCM56845 (Trident)	Supported and Tested

- Already supports Accton OCP boxes
- Working on OCP box with x86 with support from Interface Masters
- Going to work with other ODMs

Conclusion

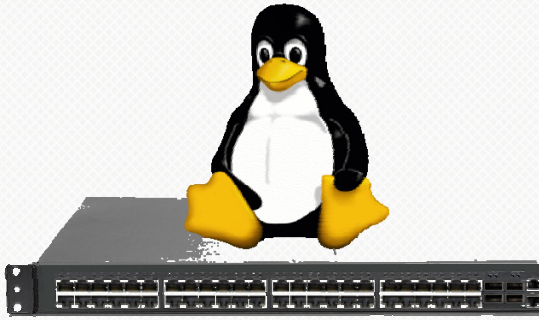

- ONL is a Linux Distribution for bare metal switches
- Growing support for OCP and non-OCP switches
- Now supports ORC forwarding agent
  - Indigo-based OpenFlow agent in progress
- Find out more at http://opennetlinux.org
  - Documentations, videos, pre-compiled binaries