

The (almost completely) Open Source Whitebox Router

ONF Connect
Session: Next-Gen SDN (P4 & Stratum & UPAN)
Santa Clara
December 6th, 2018

Hagen Woesner




Building an Open Source Router on (BCM) whiteboxes

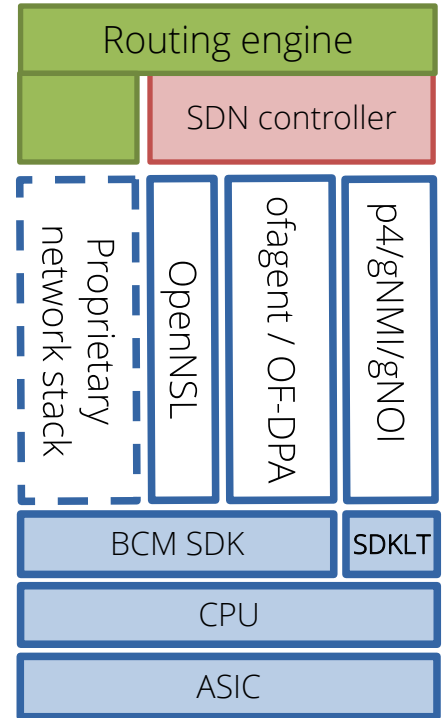
What you need:

- routing engine
- controller
- whitebox switch with OS

- ❑ all in one box
- ❑ simple configuration
- ❑ no need to write own code



- ❑ scalability (use more than one switch as one)
- ❑ ability to modify code on the router if needed,
- ❑ fix bugs
- ❑ add features



So where are we with this?

Open Source Operating System

Open Network Linux.

- Debian-based, x86-oriented
- Lots of useful work in ONLP (ONL platform specifics: LED, SFP support, fan control)
- but inherently useless for the purpose,
 - **unless** you have the .deb package with ofagent/ofdpa/OpenNSL
- Thankfully, ONF (Saurav) has put up some on the CORD pages
- Only for Accton/Edge-core switches, though



So where are we with this?

BISDN Linux

- We have developed **BISDN Linux**
 - Same, but different
 - Yocto-based
 - Linux Foundation Project, regular updates, recipes, kernels
 - ONIE image with OF-DPA 3 EA5, ofagent
 - Tested with ryu and ONOS 1.11.2, 1.12
- Running on Delta/Agema 7648, 8032 (Trident-II) 9032 (Tomahawk), 5648 (Tomahawk+), Qumran coming...
- Also: Edge-core 5712, Quanta LY8, Redstone XP

<http://repo.bisdn.de/pub/onie> and <http://repo.bisdn.de/pub/onie-testing>



SDN Controllers

- OpenDayLight and ONOS:
 - Powerful – and bulky – adapters between north-bound APIs and south-bound protocols
- Need to re-implement all network functions that are available in Linux kernel or user space
 - You want LACP, LLDP? Write an APP
 - You want Spanning Tree Protocol? Write an APP
 - You want VXLAN? OK, is it on the ASIC? Yes? Write an APP.
 - You want BGP, OSPF, IS-IS? ...
 - Ok, by now, most apps have been written, but rarely are feature-complete



Use kernel code and functions for SDN networking

Introducing baseboxd

- Idea dates back to *RouteFlow*, or even earlier...
- Baseboxd controller
 - Taps into Linux netlink, translates to OF-DPA 3.05,
 - Create tap interfaces for remote switch ports in Linux host
 - Create VLAN aware bridge
- Use systemd/networkd to persist network configuration
- grpc/protobuf and IETF/ONF yang models to the north of the **HOST**
 - grpc-ofdpa interface for non-OpenFlow calls to the switch
 - VXLAN interfaces
 - Mirror ports

<http://github.com/bisdn/basebox>



Routing engine

Quagga: extinct since 1863

- FreeRangeRouting is a recent fork of quagga
 - Very active OSS project
- Latest IETF drafts quickly implemented
- Provides
 - Zebra (kernel interaction)
 - BGP daemon
 - OSPF
 - IS-IS
 - RIP
 - Many more



Natural History Museum, Berlin

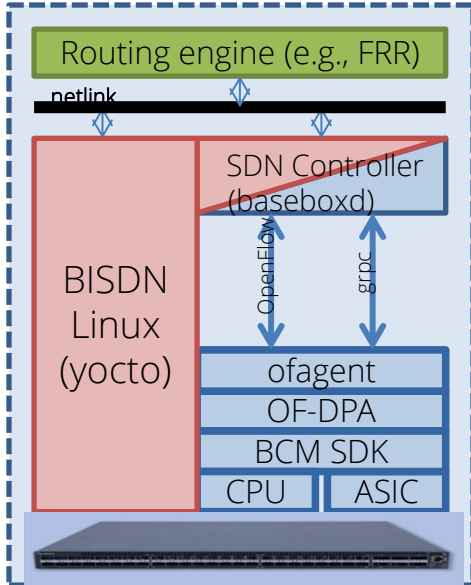
dead quagga.



There are options...

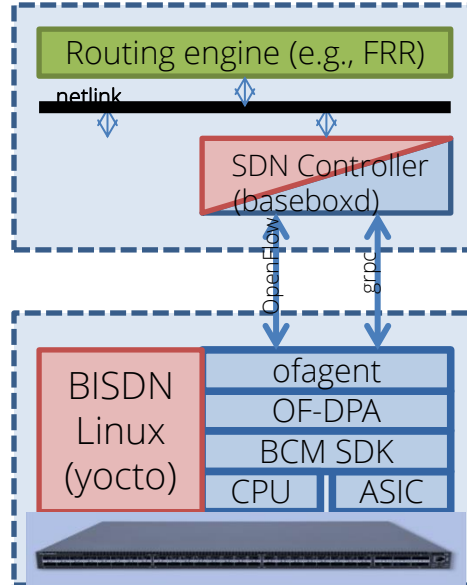
On-switch is cooler.

- Feels like a „real“ router
- Works out of the box, no initial configuration required



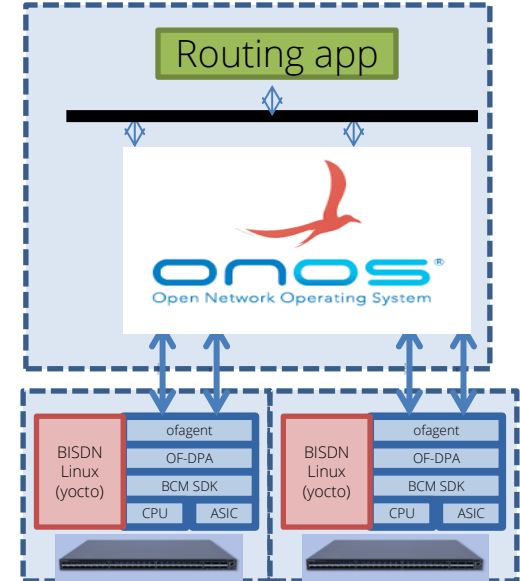
Off-switch is easier

- less kernel dependencies
- More space, infinite compute power



Centralized control scales

- Control OF-DPA switches
- Re-implement NF on top



Basic components ready and tested

BASEBOXD

- L2 (incl. source-MAC learning)
- L3 (v4 and v6 unicast)
- switch statistics via gRPC and openconfig-interfaces
- Fedora 28 systems (off-switch)
- yocto 2.5 (sumo) (on-switch)
- VXLAN

BISDN LINUX

- Support for optical SFP modules
- Scripted controller configuration (“basebox-change-config”)
- FRR as `systemd` service
- OF-DPA gRPC endpoint
- `on1pdump` support
- OF-DPA 3.0 (including group modifications)
- Mirror port (ingress/egress)



Building a Free and Open Source Router

1. Free for download ONIE image for switch

- users can still build baseboxd and FRR to run off-switch

2. Package management system

- Users can add and update packages via remote package repository

3. Expose SDK and let users build own functions

- Broadcom SDK to be provided as binary



Q: It's Open Source now, so where's your business?

A1: in Open Source

We offer support licenses and custom function development,
e.g., mirror ports

A2: The OS is commodity, we develop applications on top
e.g., BNG



Way forward: p4runtime, gNMI, gNOI



We are now members of **STRATUM**

Gradually move over from OF-DPA to p4runtime, gNMI seems close.

Contribute bug fixes for OF-DPA Community version



What we wanted to tell you...

- Climb on the shoulders of giants!
- Use existing software from living OSS projects!
- We can – together – actually fulfill the promises of Open Source Routing
- baseboxd may be the glue to make it happen. Check it out:
 - <https://github.com/bisdn/basebox>
 - <https://basebox.org>
 - https://docs.bisdn.de/setup/setup_standalone.html
 - <https://github.com/bisdn/basebox/tree/master/examples>

