

Whitebox Lab at IU

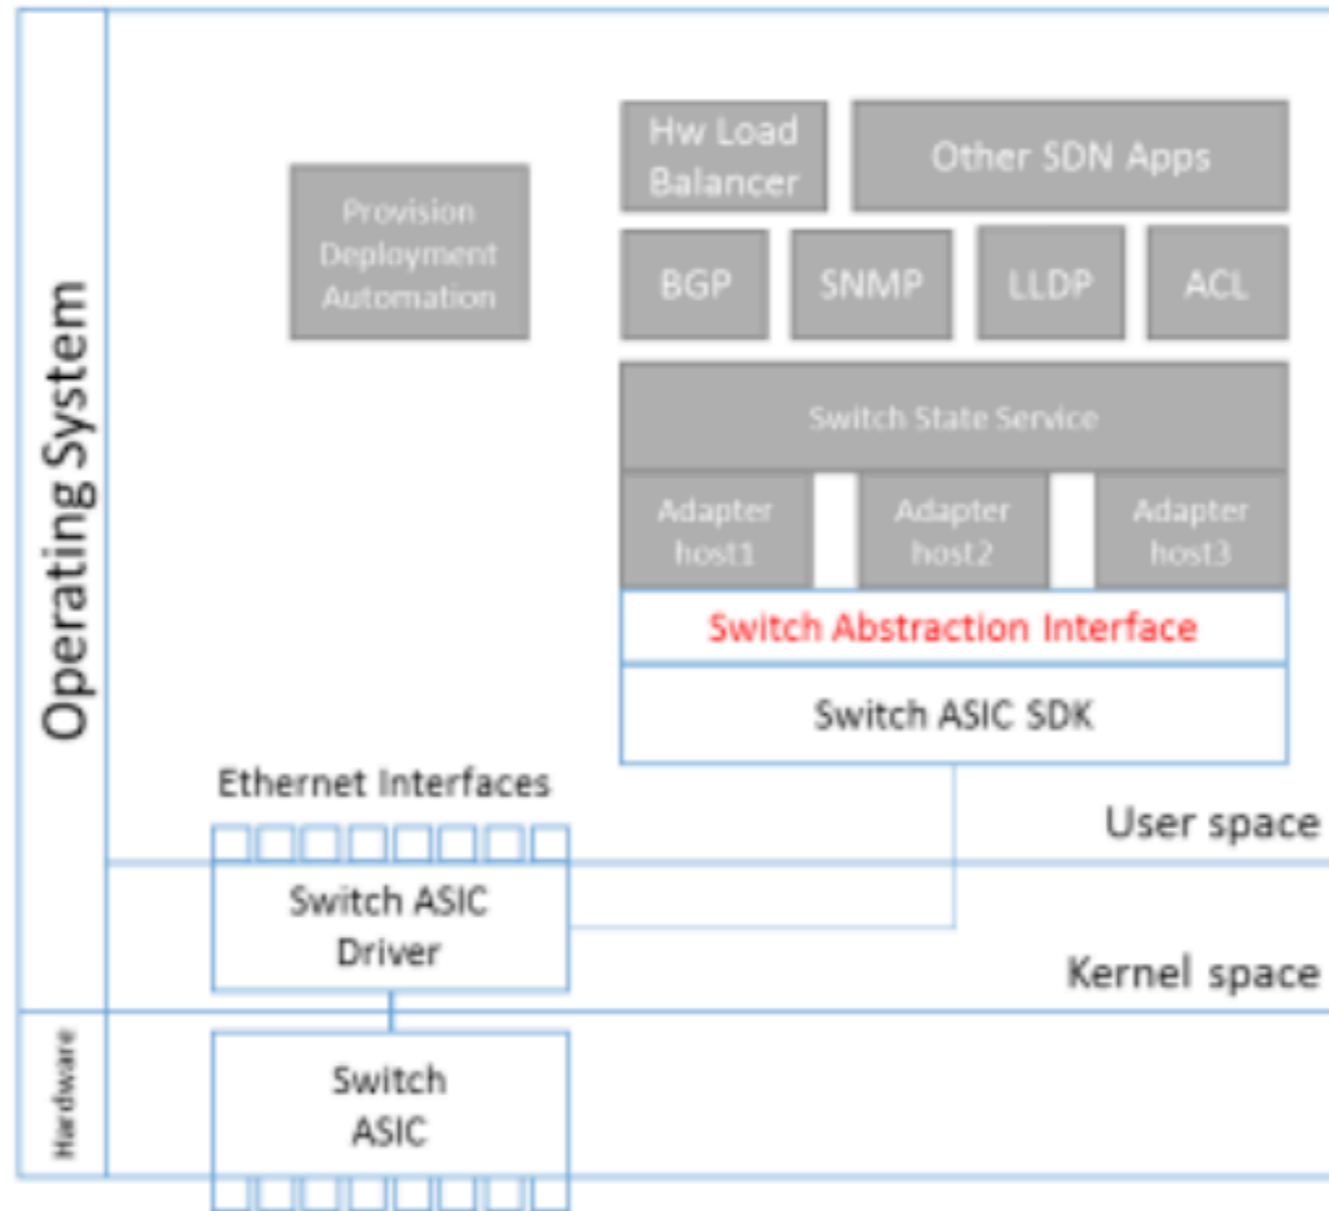
Ed Balas

Overview

- With NSF support, we began to look more closely at the emerging trend to disaggregate hardware and software known as Whitebox switching.
- This project is funded by the NSF through EAGER grant **#1535522**. We work with R&E community stakeholders and ecosystem participants to identify and evaluate use cases and network components in our lab, sharing the results of these efforts with both.
- focused mostly on campus use cases, based on available hardware at time
- more in depth coverage at
 - <https://sites.google.com/view/iu-whitebox-project/home>

Whitebox switching

- network hardware disaggregation
 - routing / switching software from hardware
- using merchant silicon
 - broadcom or other switch chipsets
 - abstraction layers to address IP concerns —> SAI
- value engineered to lower the cost per bit
 - originated in cloud providers



What this looks like

relevant projects

- Open Compute Project
 - <http://www.opencompute.org/>
- Telecom Infrastructure Project
 - <https://telecominfraproject.com/>
- Open Optical Packet Transport
 - project Voyager
 - <https://telecominfraproject.com/project/backhaul-projects/open-optical-packet-transport/>

hopes and dreams

- Use linux management paradigms to operate switches at scale
 - Ansible, Puppet, etc
- Gain access to new measurement capabilities in the new chipsets
 - directly observe buffer use —> BroadView API
- Lower cost per bit
 - disaggregated packet optical
 - only features you need for your use case

Concerns

- Is this model really supportable without your own engineering team
 - OpenFlow —> you are the system integrator concern
- What will the disaggregated software stack give us
 - will SAI be expressive enough (all features vs lowest common denominator)
 - will these things perform
- will this have relevance outside cloud providers

What we did

- Build out a lab
 - modest setup several switches, + old endpoints
- test a few things related to know use cases
 - much based on OpenFlow
- learn
- try again

Hardware

- **Edgecore: AS-6712**
 - **Ports:** 32 x QSFP+ ports, supporting 32 x 40 GbE connections or 104 x 10 GbE connections,
 - **Chipset:** Broadcom Trident2,
 - **CPU:** Intel Rangely C2538 quad-core 2.4GHz x86, 2GB DDR3
- **Agema: 7448-CU**
 - **Ports:** 48x 10GbE SFP+, 4x 40GbE QSFP
 - **Chipset:** Broadcom Trident
 - **CPU:** Freescale P2020, 4GB DDR3-SDRAM
- **Dell: S6000-ON**
 - **Ports:** 32 x QSFP+ ports, supporting 32 x 40 GbE,
 - **Chipset:** Broadcom Trident2
 - **CPU:** ?, 4GB DDR3
- **Mellanox**
 - **Ports:** 32 40/56/100GbE ports in 1RU / Up to 64 10/25/50GbE ports,
 - **Chipset:** Mellanox Spectrum
 - **CPU:** Intel x86 1.40GHZ Dual Core, 8GB DDR3 RAM, 32GB SSD

Evaluations

- SciPass Science DMZ with Mellanox switches
 - OpenFlow supported and worked for SciPass
- PerfSONAR on a switch running switchdev and fedora
 - on the Mellanox nearly 10Gbps
- Basic use and tutorials
 - ONIE - boot loader
 - OS10 - Dell NOS
 - SONIC - MS Azure NOS

Challenges

- getting use cases and community involvement
- slow going at times due to bleeding edge
- overly focused on OpenFlow in early phases
- never got to point of evaluating whether new chipset APIs can help us determine buffer requirements empirically
- multi-year time horizon for envisioned production impact

What we have learned

- Routing stacks are a work in progress
 - SoNIC and SnapRoute
- Much of hardware is optimized to data center
- Linux based approach very attractive and works
- its a crowded and evolving space
 - IPInfusion and Metaswtich
 - project voyager opens up concept more broadly
- open architecture and code a good goal
- might be time to dial back the ideological purity and look at a blended approach

Going Forward

- project winding down soon
- we have a lab with lots of gear (close to our other labs)
- growing interest in disaggregation
- still looking to collaborate and leverage this experience
 - explore capabilities and supportable models
 - smart edge / fast simple core
 - multi-tenancy