IXPs – Part of the Internet Ecosystem

- BGP control plane
- Route server
- Value-added services

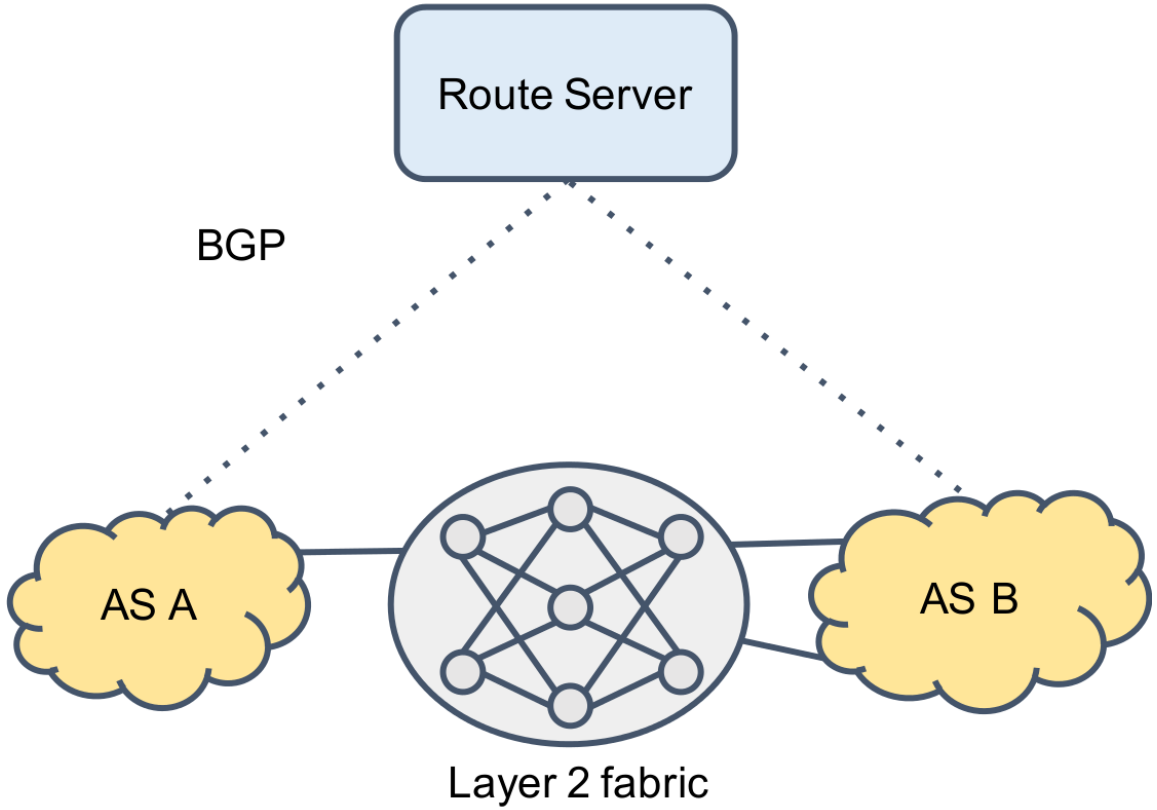

Traditional IXP

IXPs – Challenges

- Steady growth of traffic volumes
  - Disproportionately high CAPEX
- Arcane interconnection model
  - Static: long time frames, e.g. month, years
  - Myopic: limited topological visibility
  - Control-plane / reachability centric
- Limited number of value-added services

ENDEAVOUR

- Towards a flexible software-defined network ecosystem

- Goal: Bringing SDN to the Inter-domain settings

- Results are open access (e.g., GitHub [0])

ENDEAVOUR

Silver Bullet SDN?

Concrete Benefits of an Software Defined Exchange (SDX)

Members	• Port Load Balancing • Advanced Blackholing • Member Driven Monitoring
Operators	• Controlling Broadcast • Layer 2 Label Switching
Novel Services	• Port-Based Congestion Reporting • IXP as Transport Marketplace • Service Chaining

Inter-domain Networking Innovation on Steroids: Empowering IXPs with SDN

SDN USE CASES FOR SERVICE PROVIDER NETWORKS

Inter-Domain Networking Innovation on Steroids: Empowering IXPs with SDN Capabilities

Marco Chiesa, Christoph Dietzel, Gianni Antichi, Marc Bruyére, Ignacio Castro, Mitch Gusat, Thomas King, Andrew W. Moore, Thanh Dang Nguyen, Philippe Owezarski, Steve Uhlig, and Marco Canini

The authors give a comprehensive overview of use cases for SDN at IXPs, which leverage the superior vantage point of an IXP to introduce advanced features like load balancing and DDoS mitigation. They discuss the benefits of SDN solutions by analyzing real-world data from one of the largest IXPs. They also leverage insights into IXP operations to shape benefits not only for members but

ABSTRACT

While innovation in inter-domain routing has remained stagnant for over a decade, Internet exchange points (IXPs) are consolidating their role as economically advantageous interconnection points for reducing path latencies and exchanging ever increasing amounts of traffic. As such, IXPs appear as a natural place to foster network innovation and assess the benefits of SDN, a recent technological trend that has already boosted innovation within data center networks. In this article, we give a comprehensive overview of use cases for SDN at IXPs, which leverage the superior vantage point of an IXP to introduce advanced features like load balancing and DDoS mitigation. We discuss the benefits of SDN solutions by analyzing real-world data from one of the largest IXPs. We also leverage

able through more than 350 existing IXPs [1]. The largest IXPs interconnect hundreds of ASs and carry traffic volumes comparable to those of Tier 1 transit providers [2].

Over a decade of work has gone into proposing modifications to the BGP routing control plane to improve its security, and make it easier to manage and troubleshoot [3]. As these proposals require substantial global changes in BGP, unfortunately there has been no significant adoption. As such, inter-domain routing still suffers from well-known shortcomings of BGP such as its coarse-grained control of traffic based on just destination IP prefixes, and indirect control of how remote networks forward traffic.

As others have argued [4], we also deem IXPs an ideal place to spur innovation in the Internet ecosystem. First, IXPs are convergence points for a large number of ASs. Any improvement

SDX – Architecture

- Software based innovations
- Logically centralized control ease network management
- Controller designed for IXP environment

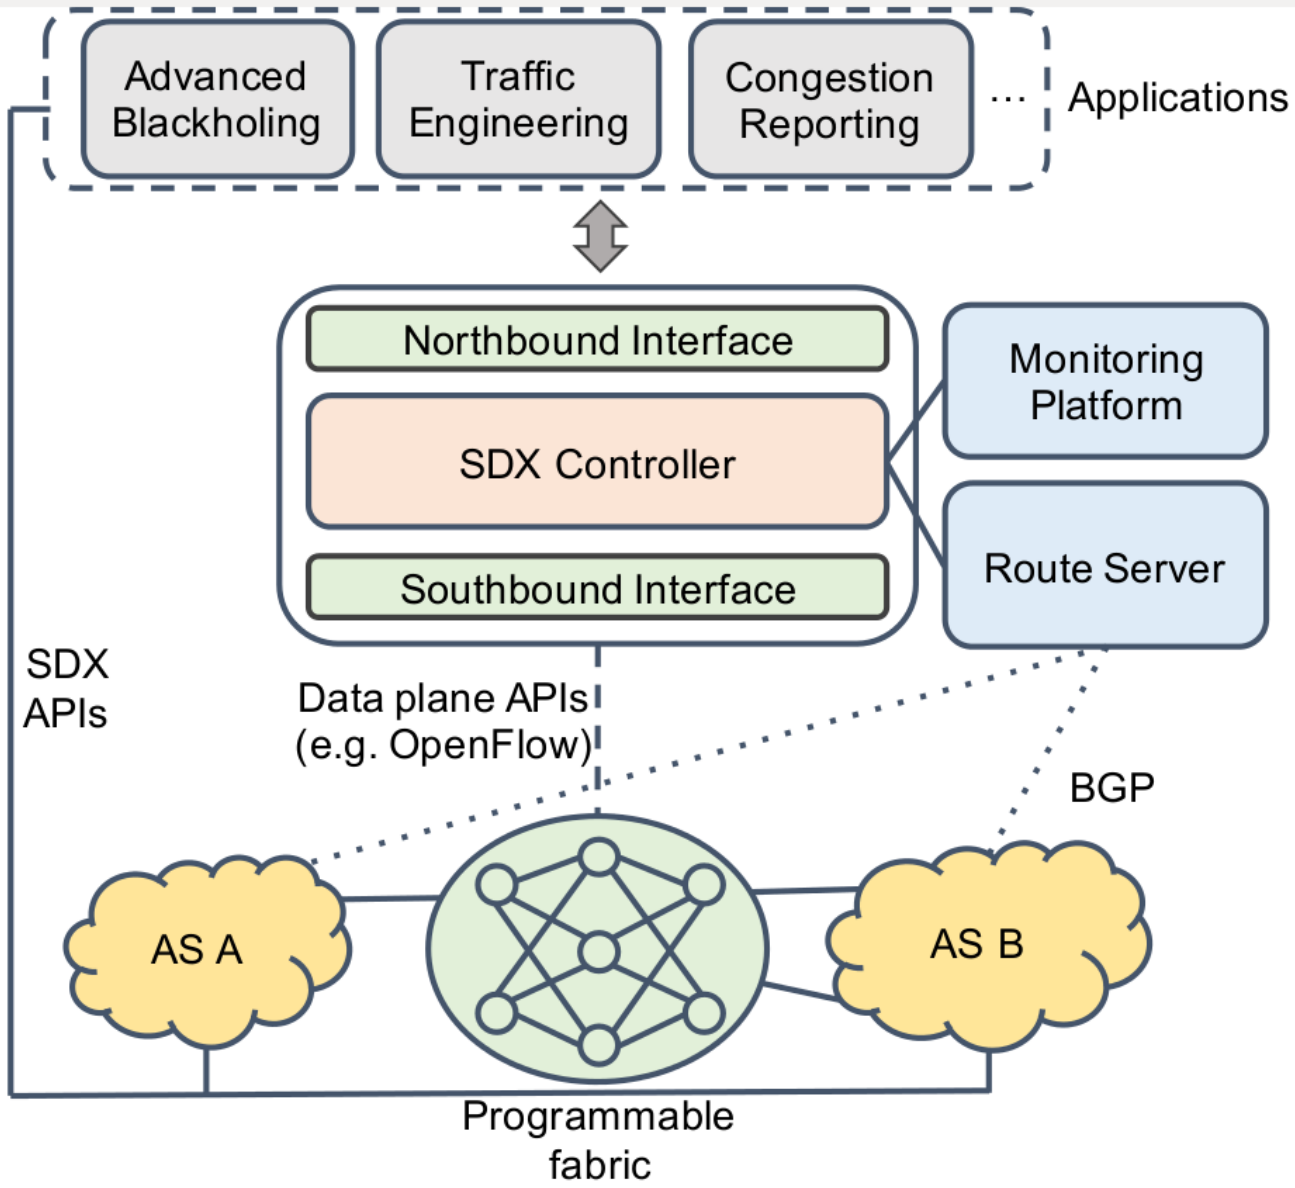

Software-Defined eXchange (SDX)

SDX Use Case – Loadbalancing

- Imperfect load balancing?
- More dynamic way to swap traffic
- Enable fine-grained routing policies

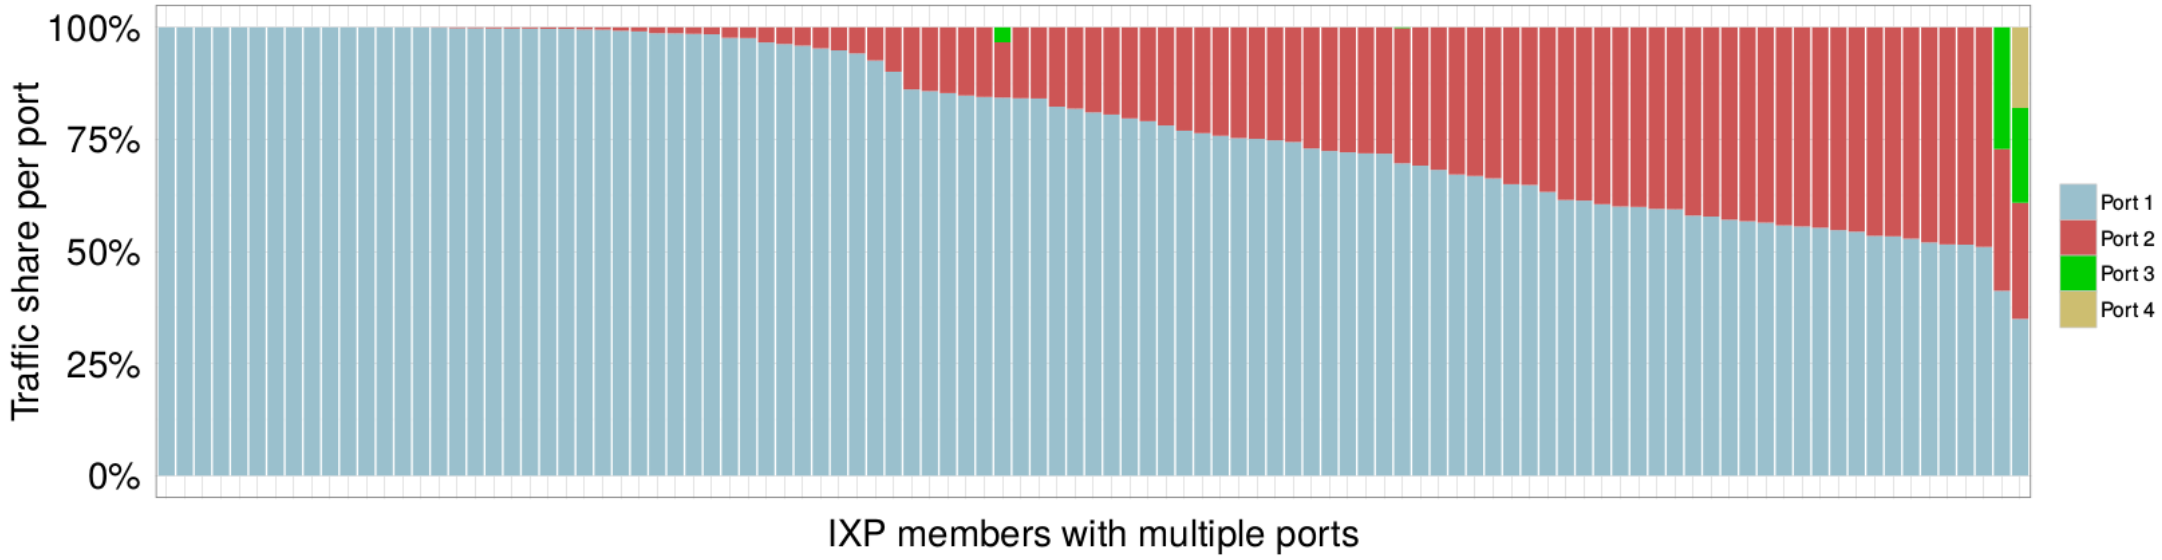

Port Load Balancing – Current Work

- Extended the SDX pipeline with a table designed for load balancing

```
cookie=0x7f, duration=227.748s, table=4, n_packets=268, n_bytes=48153,
priority=10,ip,nw_src=0.0.0.0/0.0.0.1,nw_dst=0.0.0.0/0.0.0.1 actions=write_metadata:0x30/0xffffffff,goto_table:5
cookie=0x7e, duration=227.748s, table=4, n_packets=979, n_bytes=33917,
priority=10,ip,nw_src=0.0.0.1/0.0.0.1,nw_dst=0.0.0.0/0.0.0.1 actions=write_metadata:0x20/0xffffffff,goto_table:5
```

- Basic load balancing scheme based on byte masking on IP SRC and DST
  - More sophisticated load balancing in future work

Port Load Balancing – Example

SDX Use Case – Advanced Blackholing

Today's routers can announce:

» Blackhole traffic for IP D from peer {A, all}

In an SDX future they can:

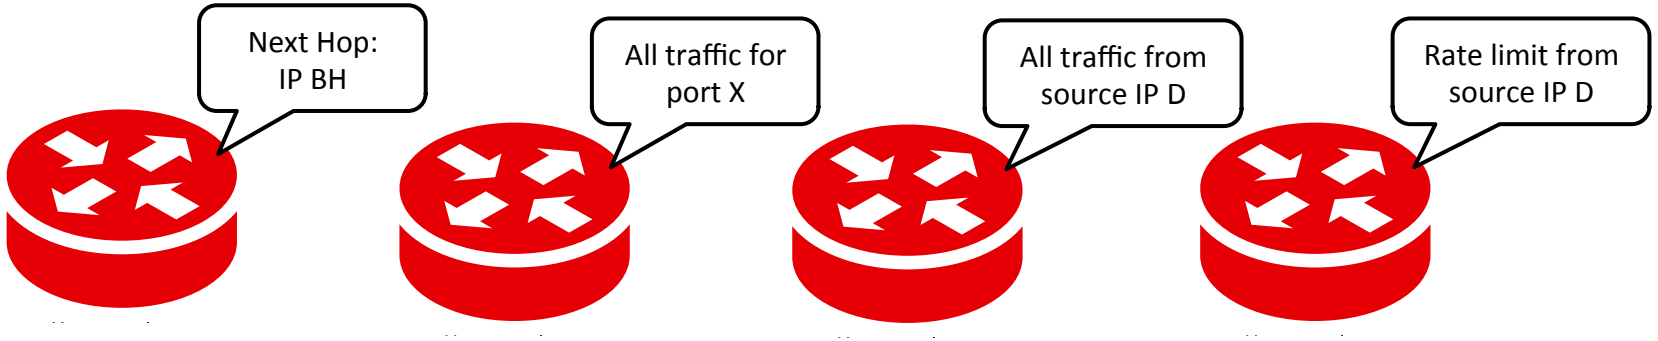

- Blackhole traffic for IP D from peer {A, all}
- Blackhole traffic for {port X, IP D, source IP D, ….} from peer {A, all}
- Rate limit traffic for {port X, IP D, source IP D, ….} from peer {A, all}

Advanced Blackholing – Current Work

- Allow the participants to install fine grained blackholing policies
  - L2: src MAC address
  - L3: src / dst IP address
  - L4: src / dst port number
- Dynamic updates of blackholing policies possible at run time
- Visualize forwarded and dropped traffic with Grafana

Advanced Blackholing – Example

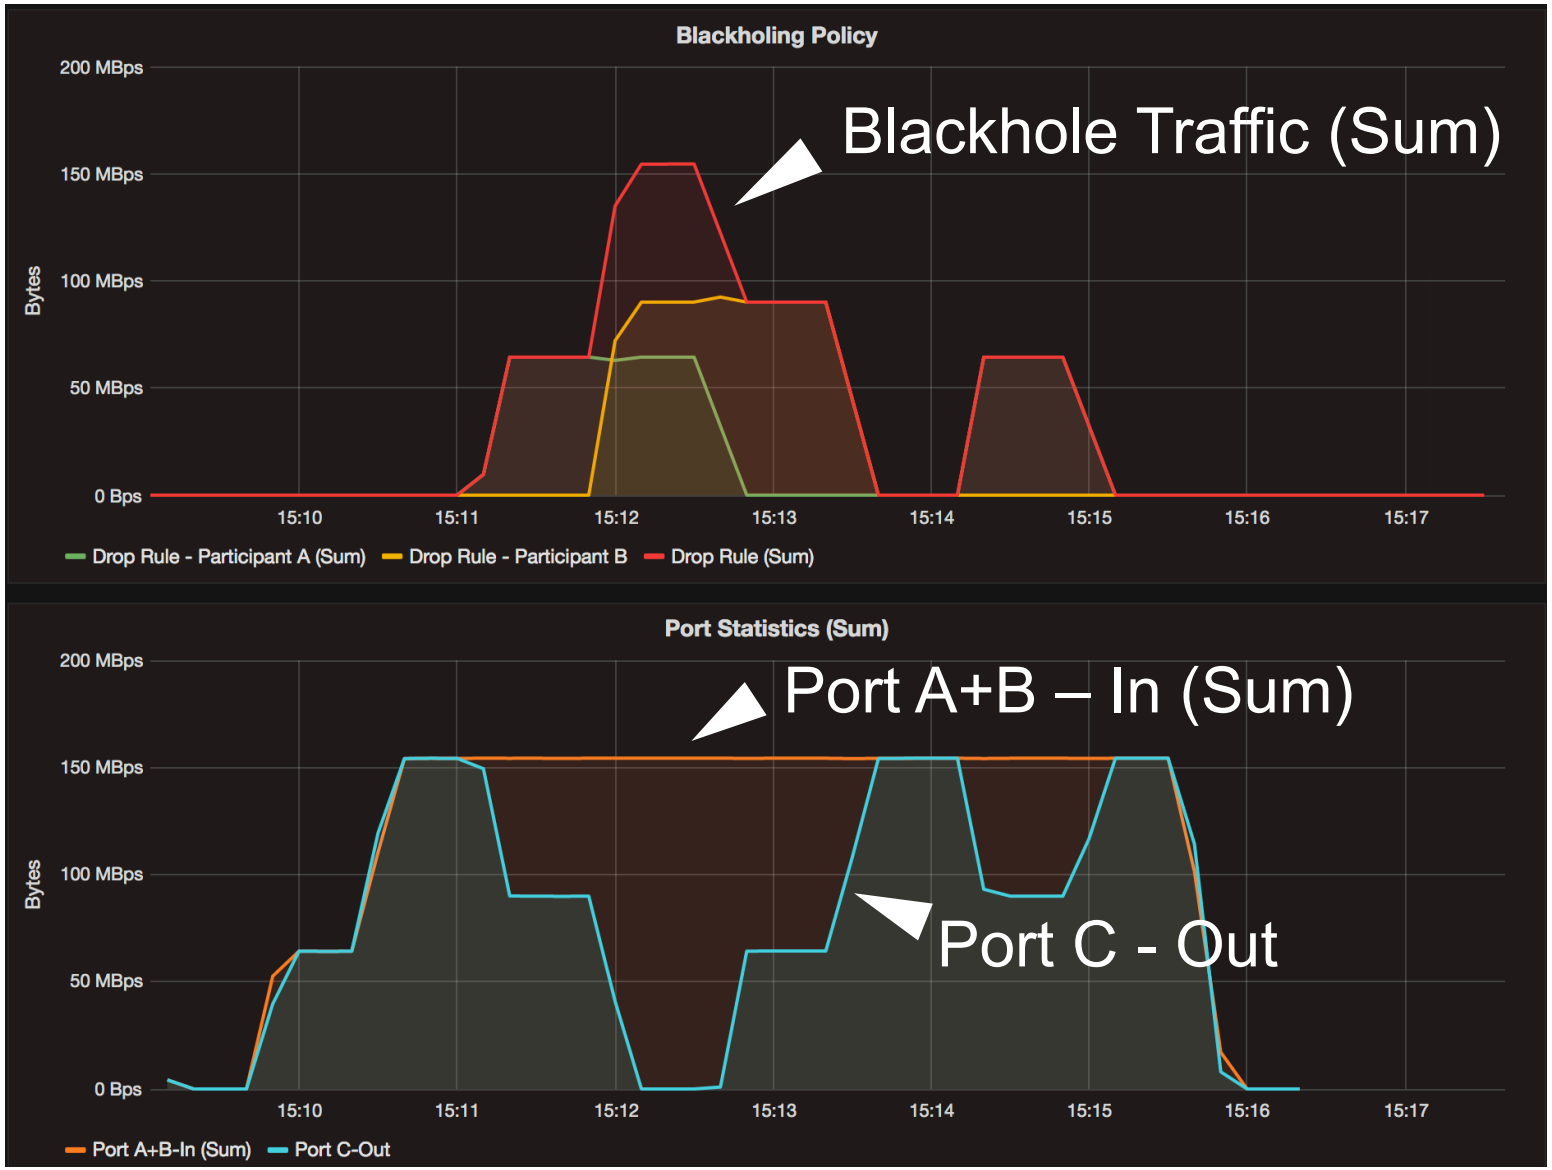

SDX Use Case – Port-based Congestion Reporting

- High imbalance of number and bandwith of ports
- Overload of egress ports
- Provide better and faster reporting for overloaded ports

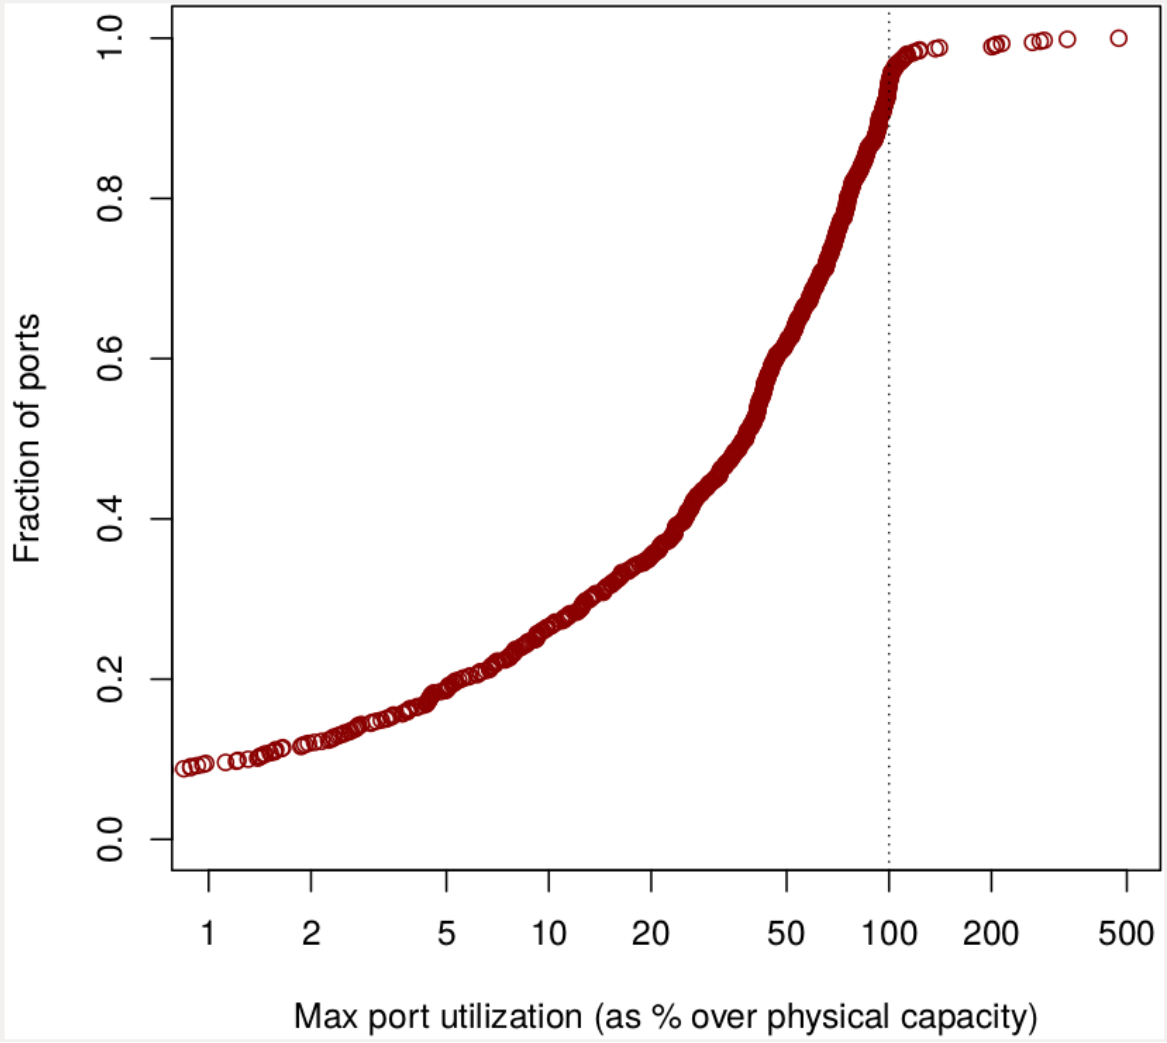

SDX Use Case – Also CAPEX Reduction?

Typical (Scale-Up) router

Scale-Out router

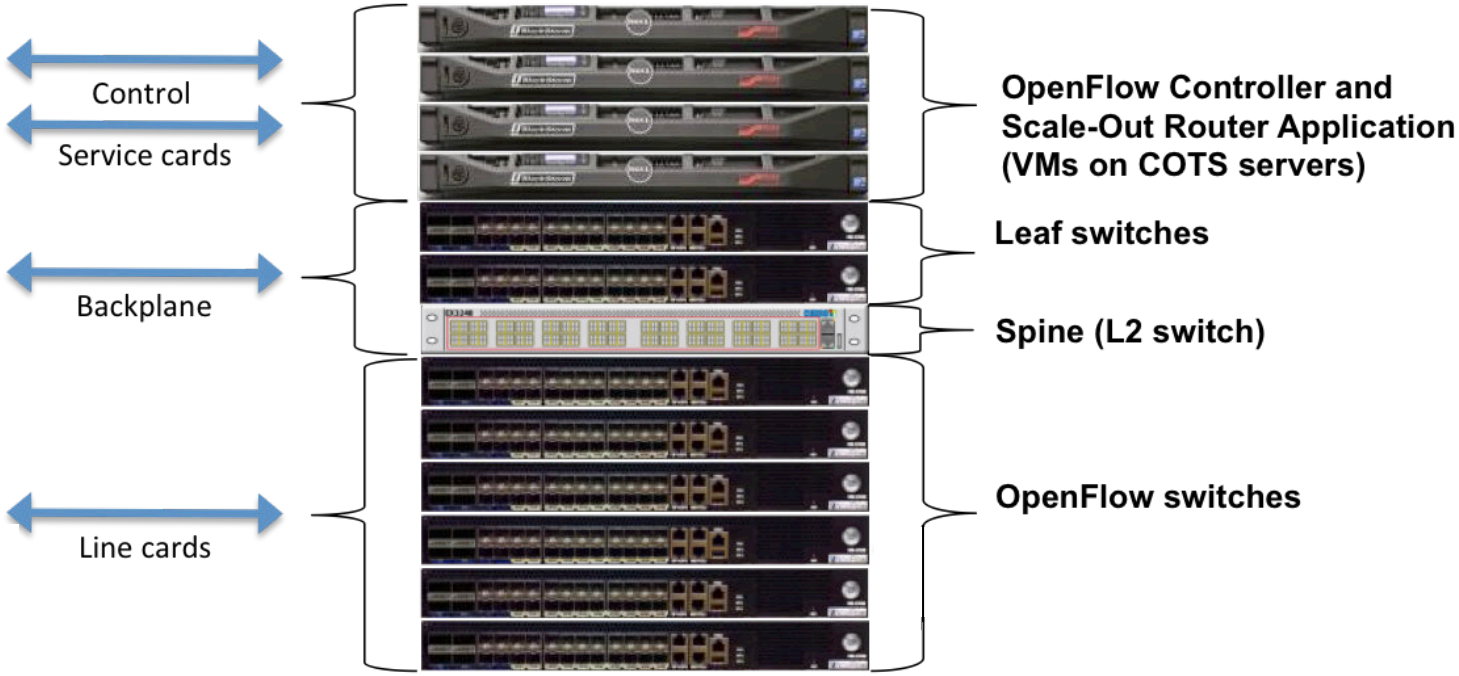

Credits: http://noviflow.com/solutions/scale-out-routers/

Limitations of SDN capable hardware

- Number on available ports (port density)
- Single table capability
  - Hard to scale with just a single table
- Limit number of TCAM entries
  - Typically in the order of 1000 – 2000
  - Up to 1 Mio. entries available leveraging DRAM
- Don‘t trust a barrier reply
  - Inconsistent states between management CPU and ASIC / Forwarding hardware
  - More complex to gather data plane forwarding state

Future SDX Challenges

1. Build reliable software 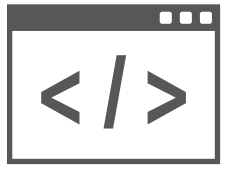

2. SDN capable hardware at scale 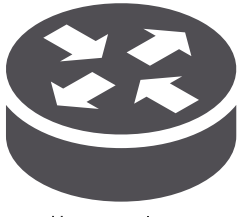

3. A lot of testing 

Credits: Code by Bernar Novalyi from the Noun Project | Beaker by Giorgio Uboldi from the Noun Project

DE CIX

http://www.h2020-endeavour.eu

http://github.com/h2020-endeavour/endeavour

E-Mail: rnd <AT> de-cix.net

Where networks meet

www.de-cix.net