



ESnet

ENERGY SCIENCES NETWORK

Virtualized Network Control: SDN Research

Inder Monga

Division Deputy, Scientific Net. Div.

CTO, ESnet

DOE NGNS PI Meeting

Washington DC, September 17th, 2014




U.S. DEPARTMENT OF
ENERGY
Office of Science



Goal of SDN Research

- Focus on ESnet use-cases for the next-generation network, ESnet6
 - Multi-layer, multi-domain, multi-vendor network
- Target science use-cases and applications especially large science flows (elephant flows)
- Influence the larger SDN industry to be aware of the science requirements
 - Enhance OpenFlow to meet those requirements

SDN Investigations for this project




ONS = Open Networking Summit
SDN = Software Defined Networking

**Network
Virtualization
'One Switch'
SC12, Nov 2012**



New Network Abstraction: “WAN Virtual Switch”

WAN Virtual Switch



Simple, Multipoint Services, Programmable

Configuration abstraction:

- Expresses desired behavior
- Hides implementation on physical infrastructure

It is not only a concept, but prototype
Demonstration with NEC and Ciena


“OpenFlow Programmable” by end-sites

Program flows:

Science Flow1: $A \rightarrow B$, QoS, Label

Science Flow2: $A \rightarrow C$, VLAN

Science Flow3: $A \rightarrow B, C$



**World's first
Transport SDN,
Dec 2012**



OpenFlow paradigm extended to L1/ λ

- **No uniform way to manage Multi-vendor, Multi-layer networks or multi-domain optical bypass**
 - Packet and Optical devices managed in totally different ways
 - Multi-vendor networks are islands
 - No global topological visibility
 - GMPLS and vendor-specific hacks needed in current approach

Packet World

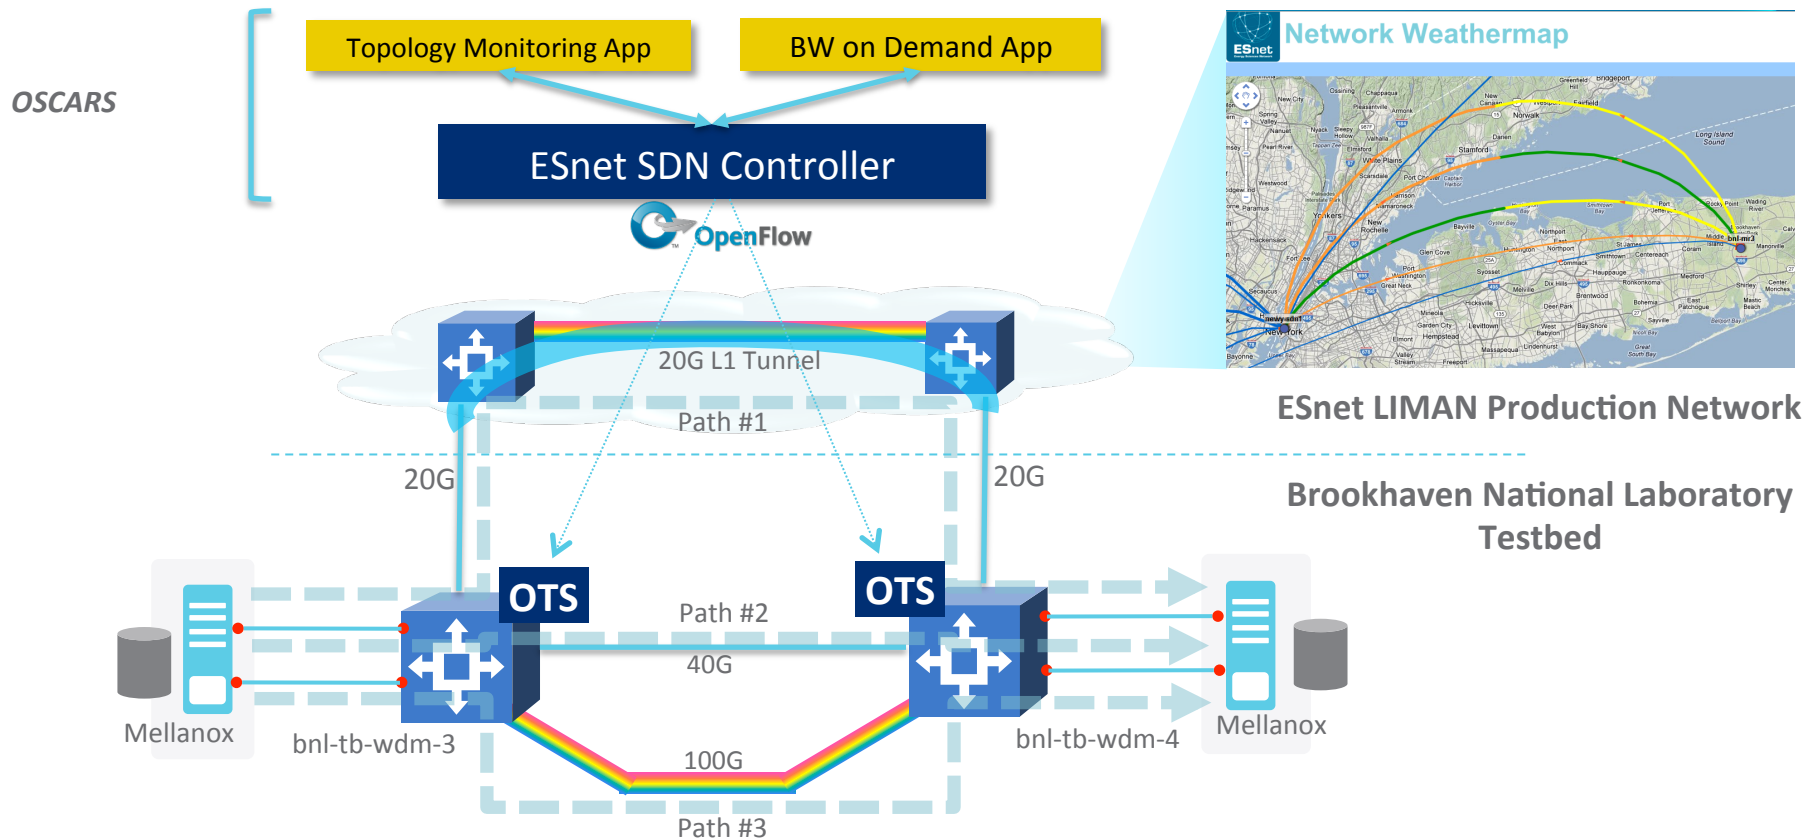
- Connectionless
- Enterprise origins
- Dynamic flows
- Inband control plane
- Numerous distributed CP solutions
- Monolithic, closed systems

Transport World

- Connection (circuit) oriented
- Service provider origins
- Static pipes
- EMS/NMS + Cross-connect paradigm
- Nascent CP (GMPLS)
- Open, programmable systems

Source: Infinera


Transport SDN Demonstration



- SDN Controller communicating with OTS via OpenFlow extensions
- Bandwidth on Demand application for Big Data RDMA transport
- 3 physical transport path options (with varying latencies)
- Implicit & explicit provisioning of 10GbE/40GbE services demonstrated



**SDX REANNZ/
ESnet, ONS,
March 2013**



SDX= Software-Defined Exchange



Demonstrated SDN-Router @ Open Network Summit, March 2013

Front-Line Assembly

DEMO

First international BGP peering using SDN in production between two national-scale network providers

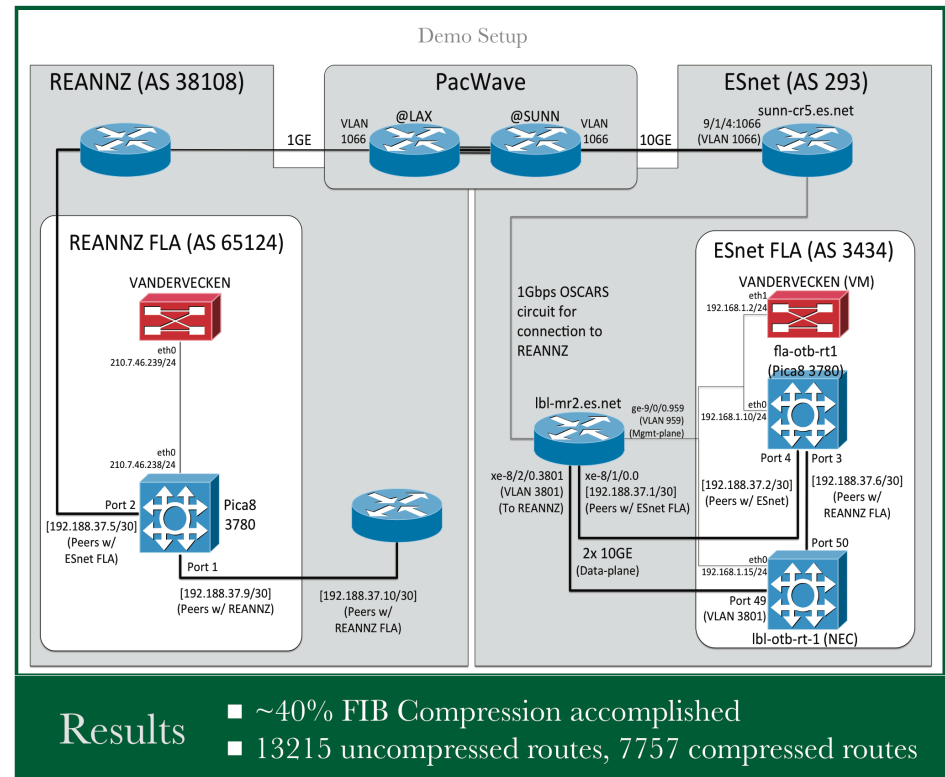
- Innovative FIB compression enables using commodity OpenFlow switches for peering
- Leverages community open-source packages. RouteFlow and Quagga
- Collaboration with REANNZ and Google
- Live SDN BGP peering across the Pacific!

Demonstration Team:

Google Network Research – Josh Bailey, Scott Whyte
REANNZ – Dylan Hall, Sam Russell, James Wix, Steve Cotter
ESnet – Inder Monga, Chin Guok, Eric Pouyoul, Brian Tierney
Acknowledgements - Joe Stringer

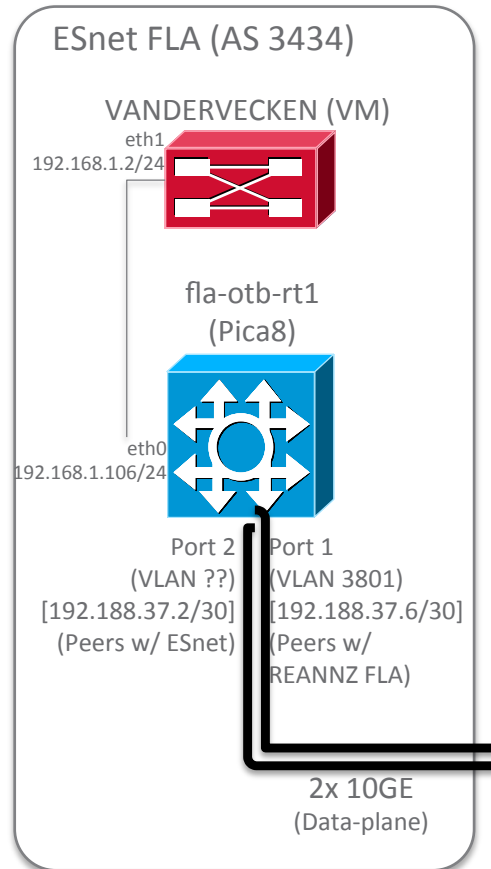


NGNS PI meeting



Software-Defined Exchange (SDX)

March 2013



Treehouse Project

(REANNZ/ESnet/Google)

12

18/09/14

April 2013

SDX: A Software Defined Internet Exchange

Nick Feamster[†], Jennifer Rexford^{*}, Scott Shenker[‡], Dave Levin[◊], Russ Clark[†], Josh Bailey^{*}

^{*} Google, [†] Georgia Tech, [◊] University of Maryland, ^{*} Princeton University, [‡] UC Berkeley

December 2013

Operationalizing SDN: A multi-agency program review

August 2014

SDX: A Software Defined Internet Exchange

Arpit Gupta[†], Laurent Vanbever^{*}, Muhammad Shahbaz[†]
Sean P. Donovan[†], Brandon Schlinder[‡], Nick Feamster[†], Jennifer Rexford^{*}
Scott Shenker[◊], Russ Clark[†], Ethan Katz-Bassett[‡]

[†]Georgia Tech ^{*}Princeton University [◊]UC Berkeley [‡]Univ. of Southern California



NGNS PI meeting

What is a 'Software-Defined Exchange'?





Image from SDX NSF workshop report

Washington DC, June 2014

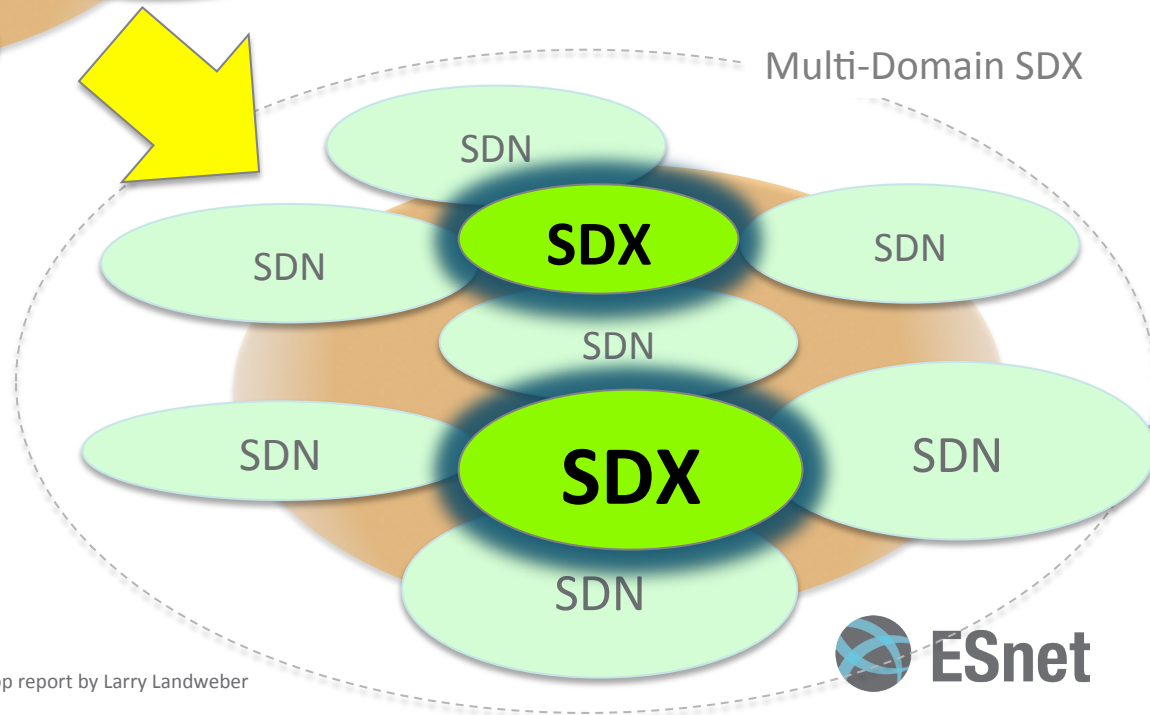
- SDX =
 - New way to build an Internet Exchange Point (IXP)
 - New way to build a Layer 2 GOLE
 - New approach for multi-domain services
 - New approach to SDI/Network Functions Virtualization
- Σ SDX = New way of building software-defined networks?

Multi-domain SDXs*

Today: “SDN islands”
GENI slices & VLAN stitching
help point the way




Next Step: Add SDX's
Build a “Rev 0” control plane,
run native next-gen apps
and scientific instruments
spanning multi-domain SDNs



Multilayer Transport SDN, October 2013



The hidden multi-layered infrastructure

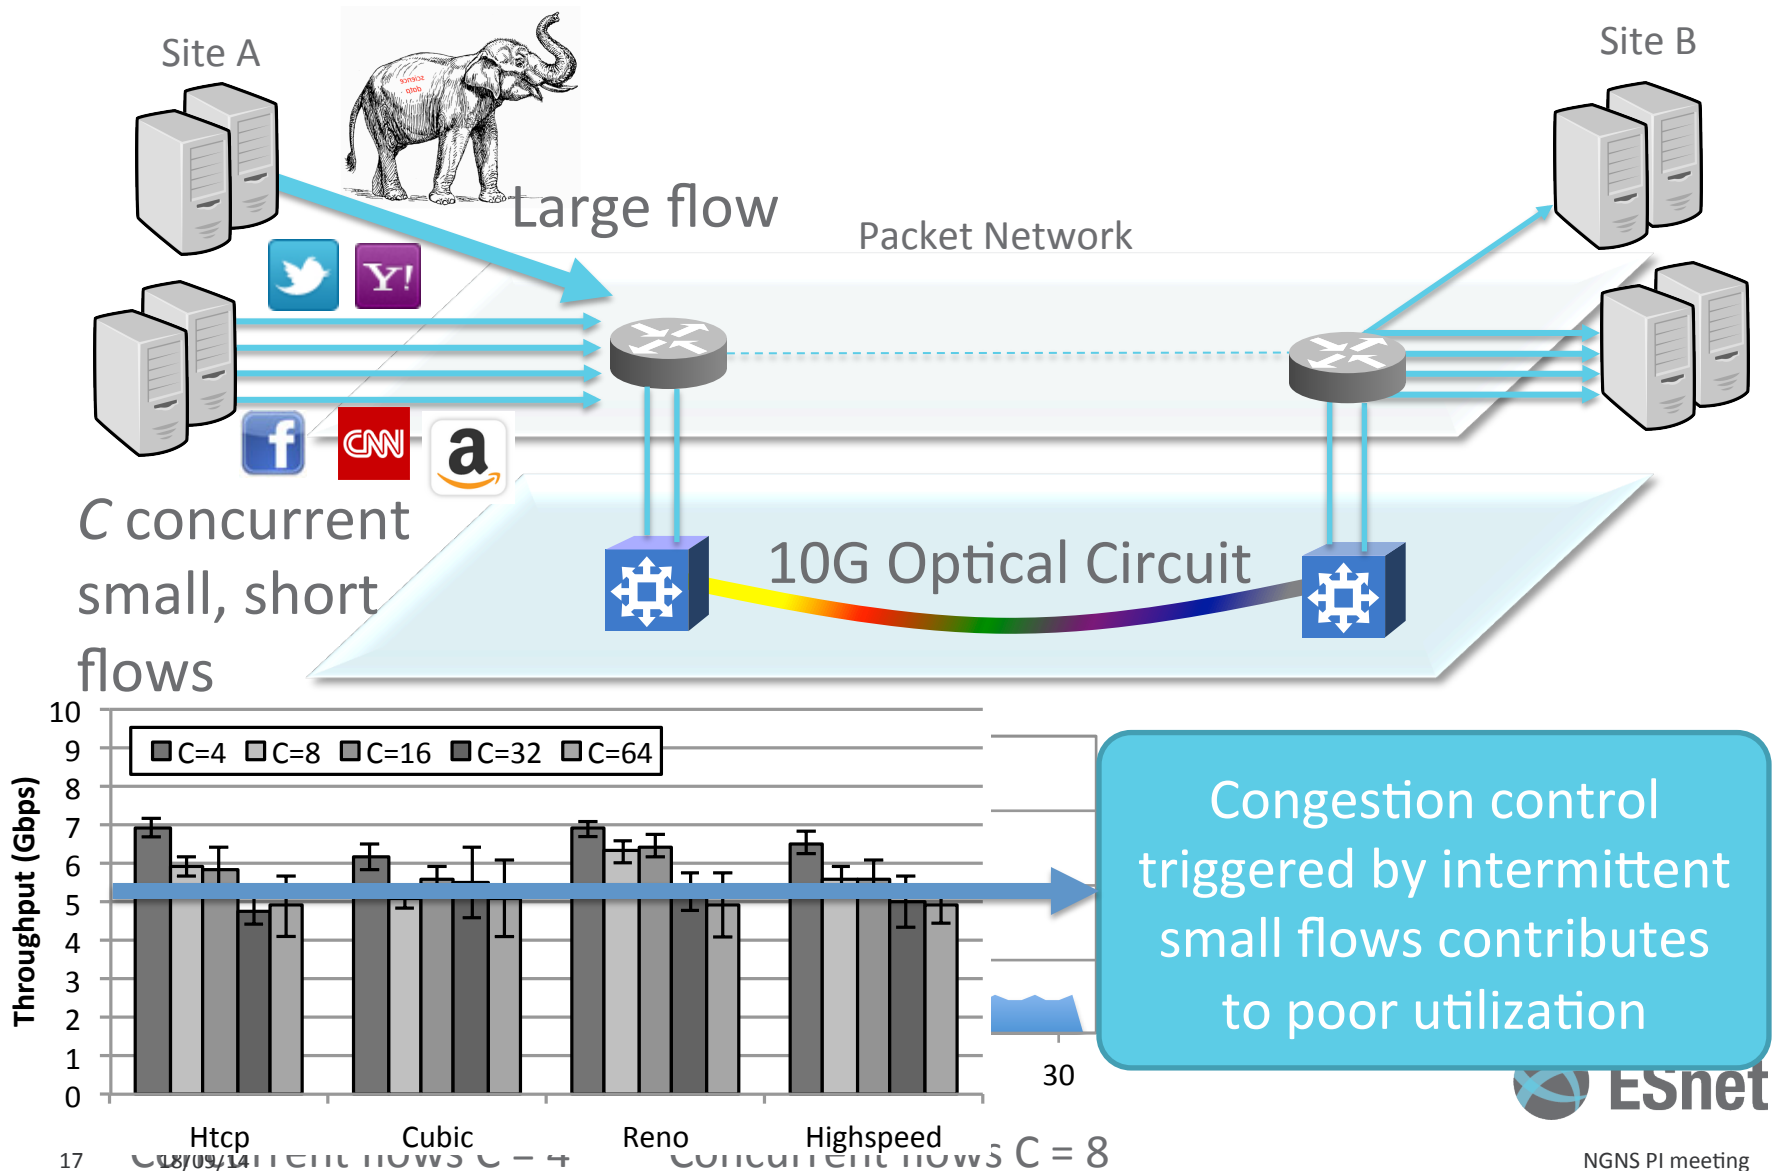


GMPLS, TL1

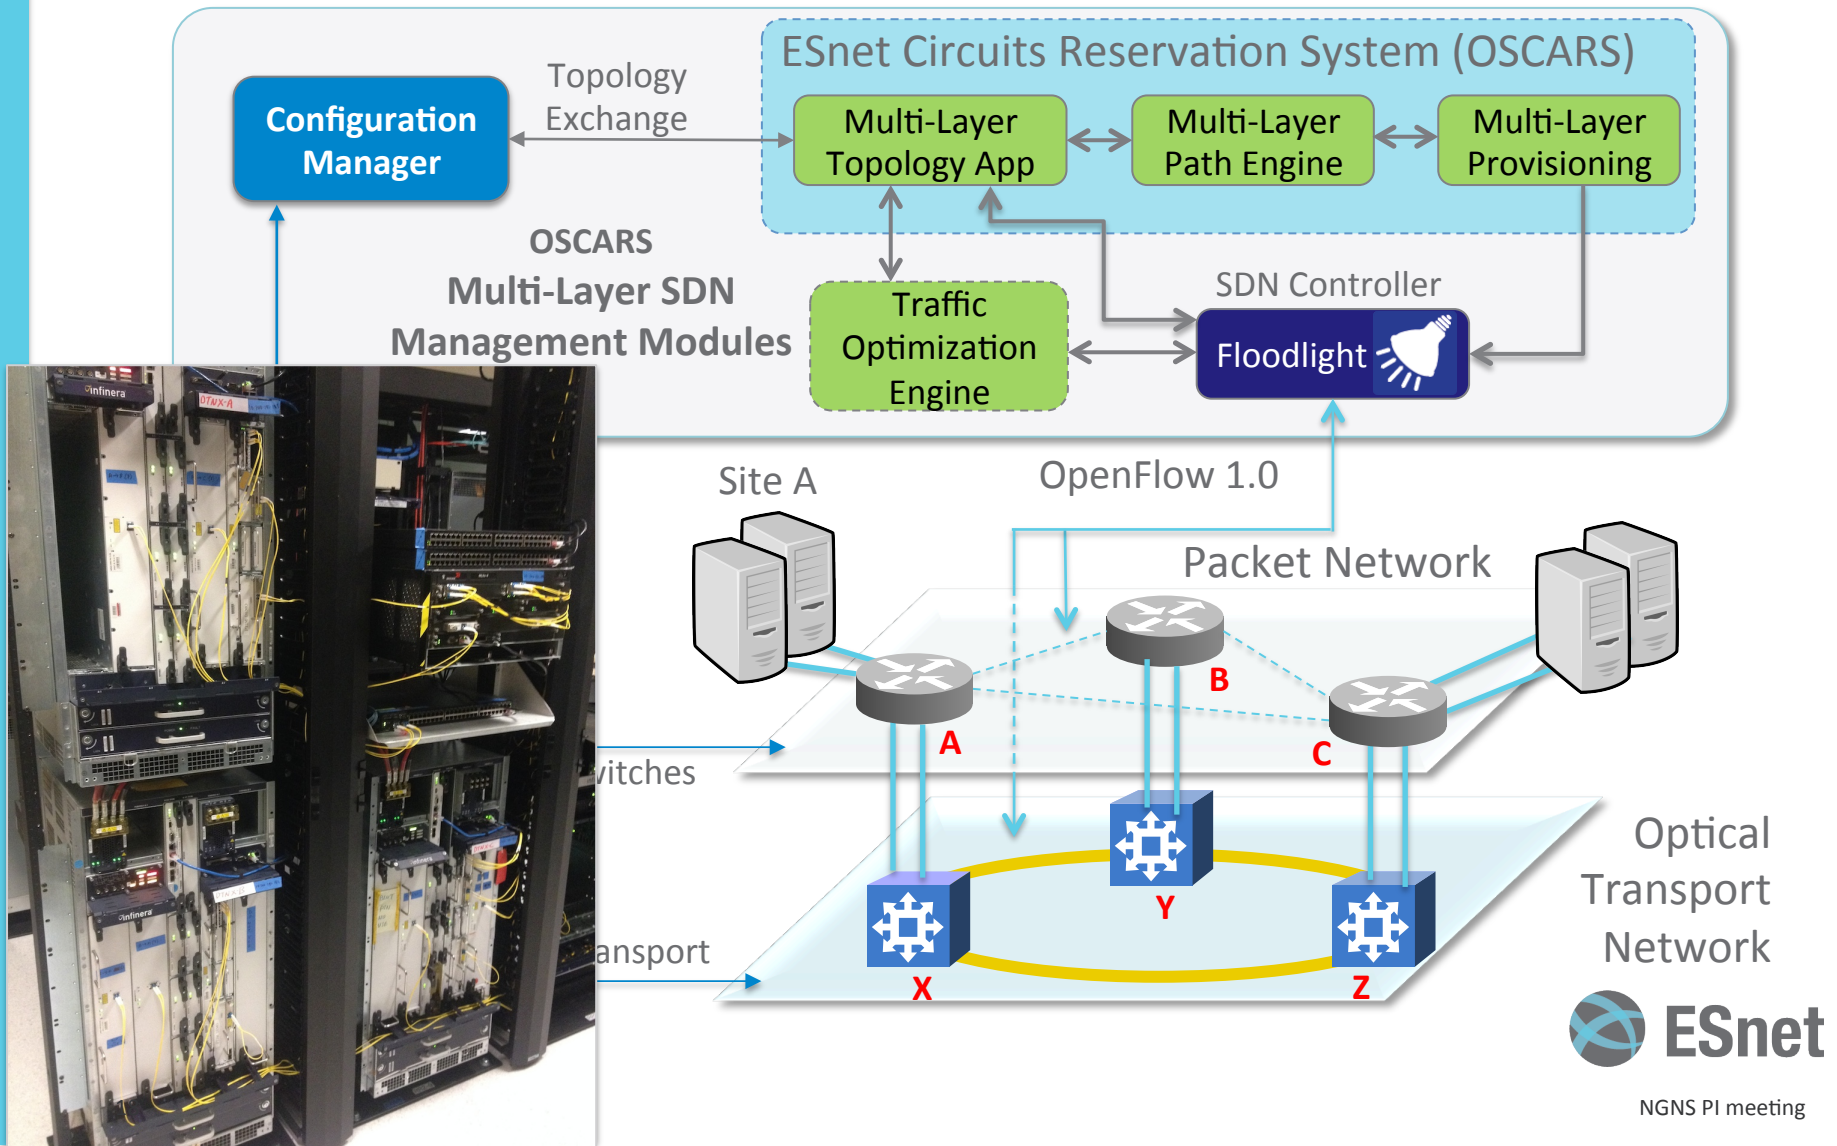
Can we manage all layers
using a unified abstraction?

Proprietary
Manual Operation

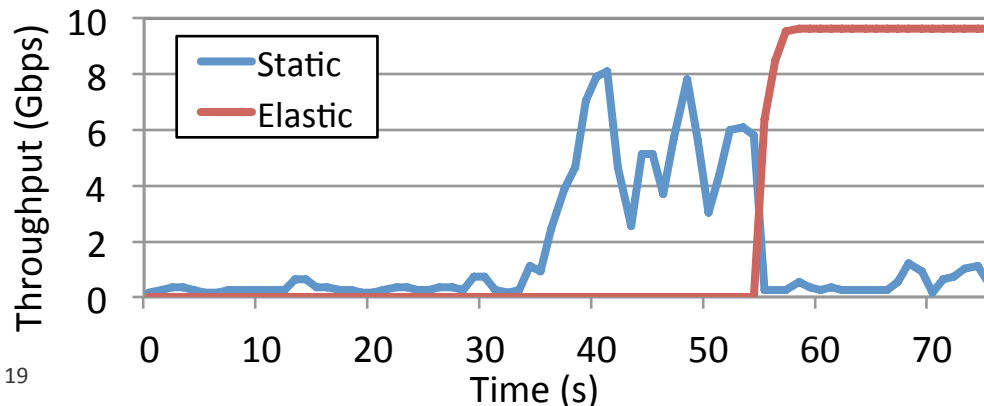

Distinct TCP Flows vs. Utilization



Multi-layer SDN



Enabling predictable application performance



T = 0: Only small flows
T = 30: Large data transfer started
T = 55: Large data transfer offloaded to dynamically allocated circuit

Funding and Research Impact

Best Student Paper Award

- Papers published
 - “Traffic Optimization in Multi-Layered WANs using SDN”, Hot Interconnects 2014
 - H. Rodrigues*, I. Monga†, A. Sadasivarao, S. Syed, C. Guok†, E. Pouyoul†, C. Liou and T. Rosing*
 - “Open Transport Switch - A Software Defined Networking Architecture for Transport Networks”, ACM SIGCOMM HotSDN Workshop, August 17, 2013
 - Abhinava Sadasivarao, Sharfuddin Syed, Chris Liou, Ping Pan, Andrew Lake, Chin Guok, Inder Monga
 - “Software Defined Networking for big-data science”, Super Computing 2012, November 11, 2012
 - Inder Monga, Eric Pouyoul, Chin Guok
 - Paper and presentation at SC
- Many talks at conferences
 - OFC, OnVector, Globecom, HotSDN, Hot Interconnects, NSF workshops...

*: University of California, San Diego, †: ESnet, Rest: Infinera



Funding and Research Impact

- Industry and R&E collaborations
 - Ciena and NEC
 - Network Virtualization demonstration
 - Infinera
 - Worlds' first Transport SDN demonstration
 - REANNZ/Google Network Research (and now SANREN)
 - Software-Defined Exchanges, SDN-based router
 - Brocade and Infinera
 - Multi-Layer SDN

Funding and Research Impact

- Industry Impact
 - Optical Transport Working group ratified in ONF after demonstration, that resolved a lot of open questions on need and viability of this approach
 - Nominated by Industry participants to be “ONF Research Associate” – three companies sponsored (independently of each other)
 - Software-Defined Exchanges (SDX), an important recommendation of the December ‘Operationalizing SDN’ report and part of NSF’s IRNC proposal
 - Multi-layer SDN demo featured at SDNCentral as their Demo Friday

What questions does our research motivate us to ask?

- End-to-end remains difficult because service guarantees have to be consistent across all the domains- describing policy (and dissemination)
- Lack of a good language to describe complex policy
 - E.g N can only use X b/w on links 1-3 with specific on-peak/off-peak rates on weekdays)
- Negotiation of resources, especially multi-domain can be non-trivial due to scope and potentially large number of variables
- What are the right abstractions other than a point-to-point [virtual] link that science applications will find useful?
- Operationalizing SDN – what are the basic manageability features needed before SDN can become production ready?