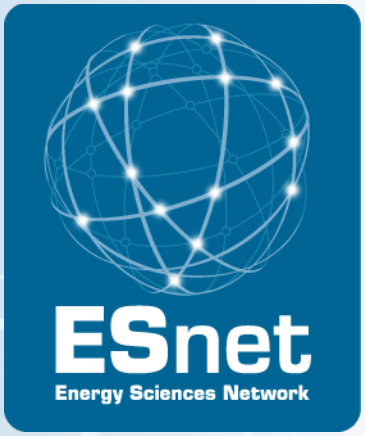

Software Defined Networking @ESnet

Inder Monga

Chin Guok

OASCR SDN Workshop May 12th 2014

Goals

- Highlight ESnet efforts in SDN
- Discuss future directions for R&E community

Funding background

OASCR Research Funding: 1 year funding to investigate SDN for ESnet

Program Funding: Continue the research and prototyping work for SDN

SDN Investigations @ ESnet brief timeline

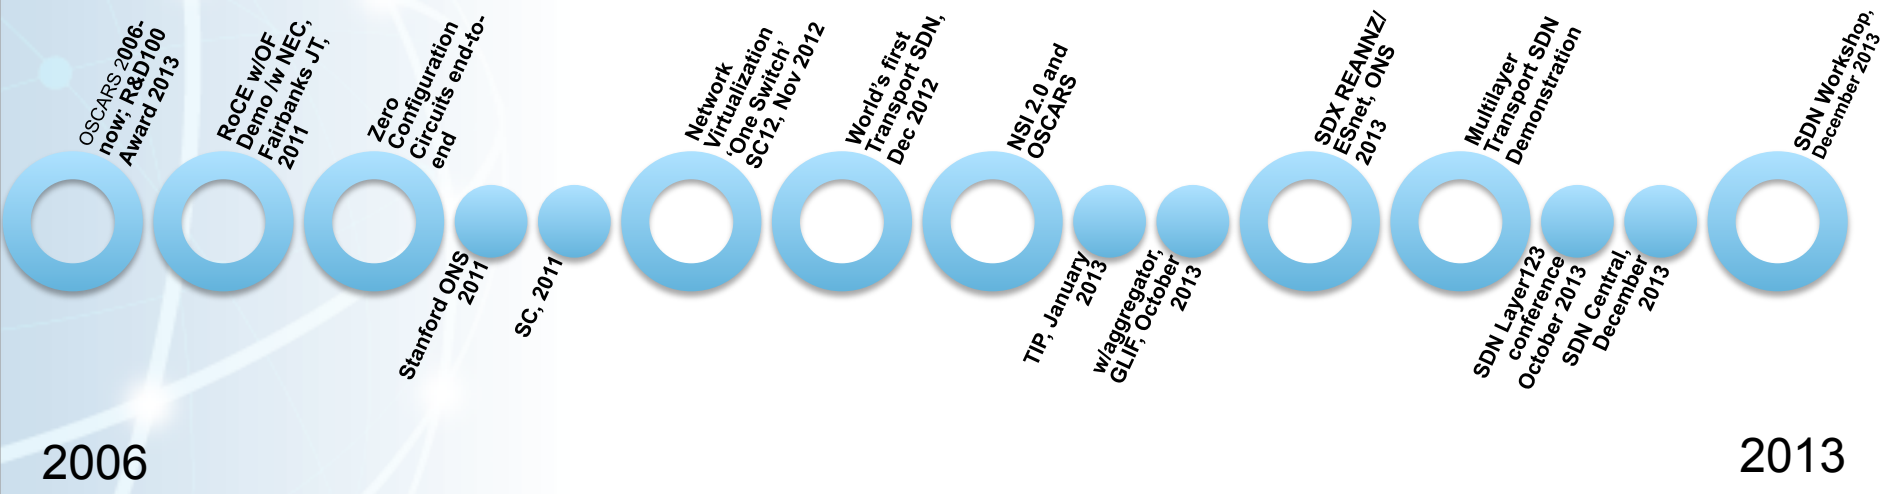

2006

2013

ONS = Open Networking Summit
TIP = Joint Techs with Univ. Hawaii
RoCE= Remote DMA over Converged Ethernet
GLIF = Global Lambda Integrated Facility
SDN = Software Defined Networking

Journey with Programmability

Joint Techs Summer 2011, Fairbanks, Alaska

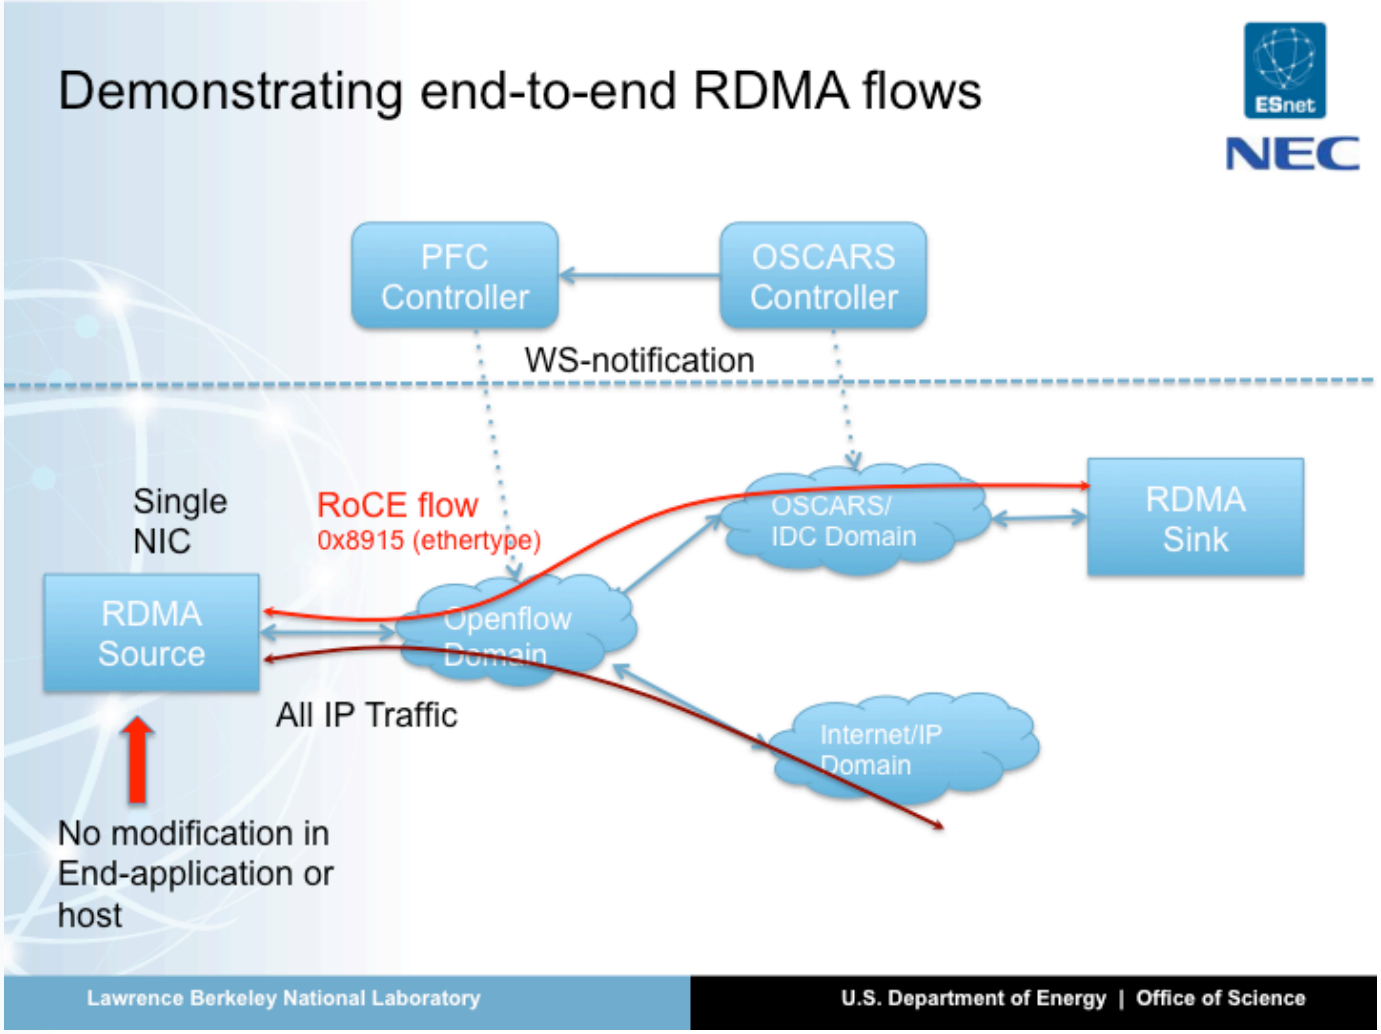

Insights

- SDN not immune from end-to-end problem
- 'unmodified end host' an attractive architecture

Microsoft ONS 2014: Implemented our first idea in production

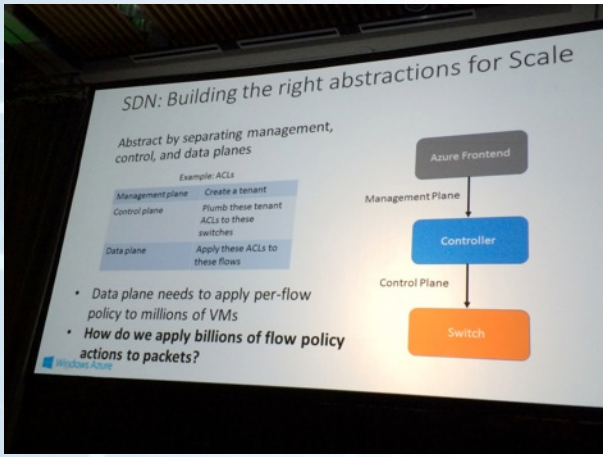

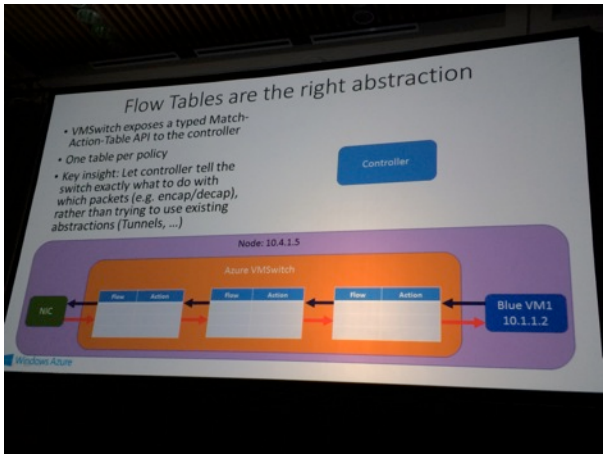

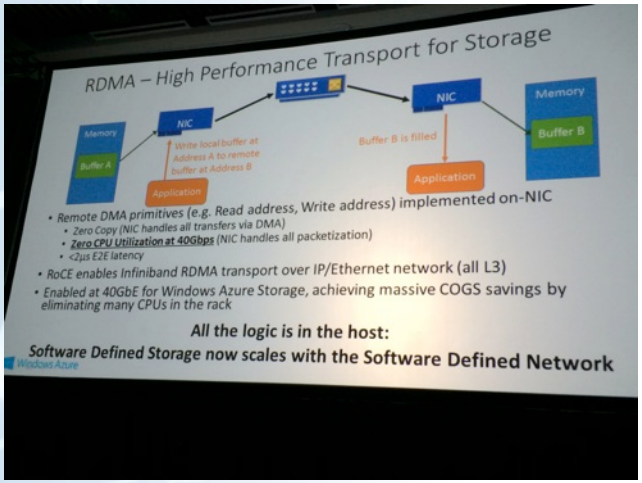

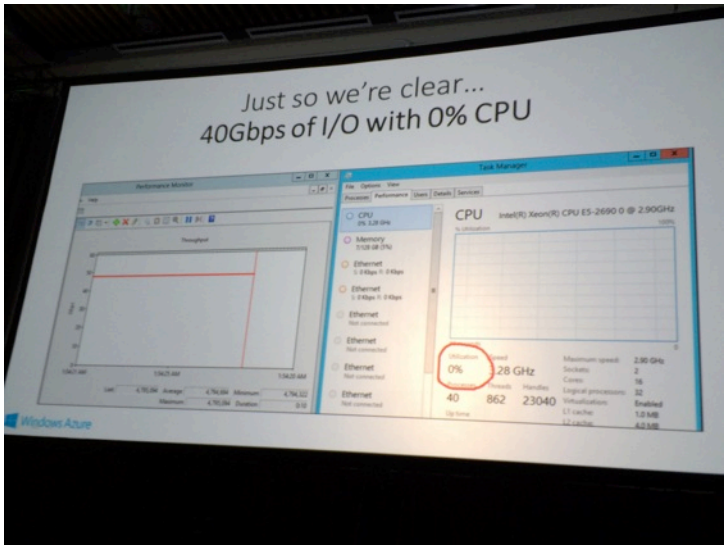

Journey with Programmability

Inaugural Open Network Summit, 2011 (Stanford) and SC 2011 (Seattle)

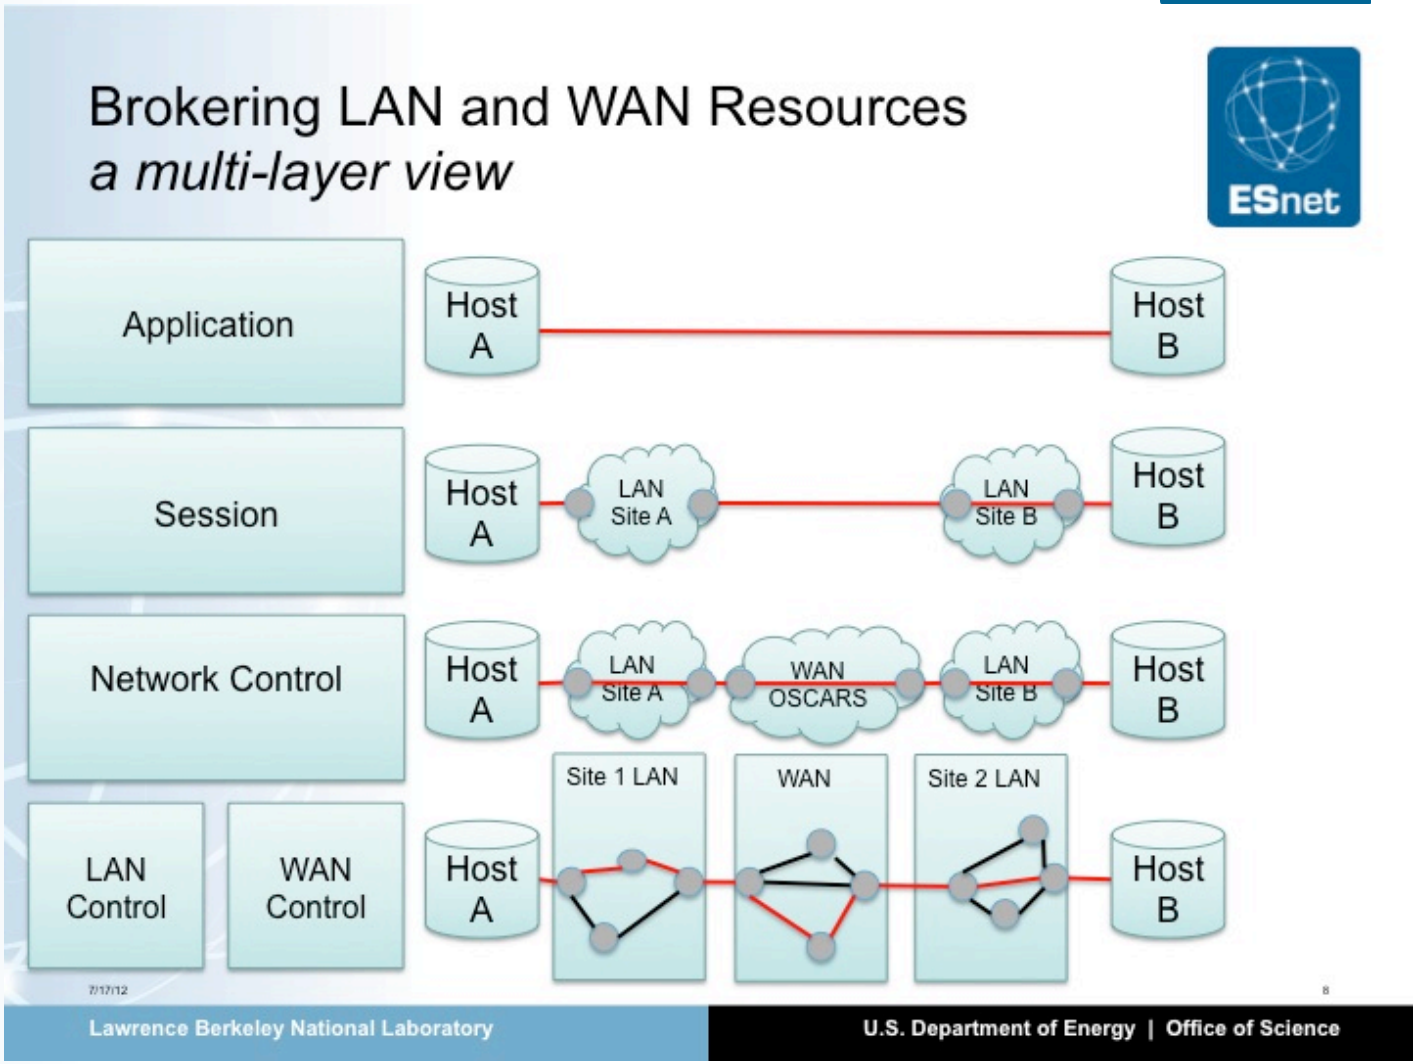

Insights

- **Zero configuration circuit**
- end-to-end communication at each layer important

(even at layer 8)

Journey with Programmability

SRS, Ciena, SuperComputing 2012, Salt Lake City

Insights

- **Network virtualization is SDN 'killer app'**
- 'virtual switch' abstraction in the WAN holds promise

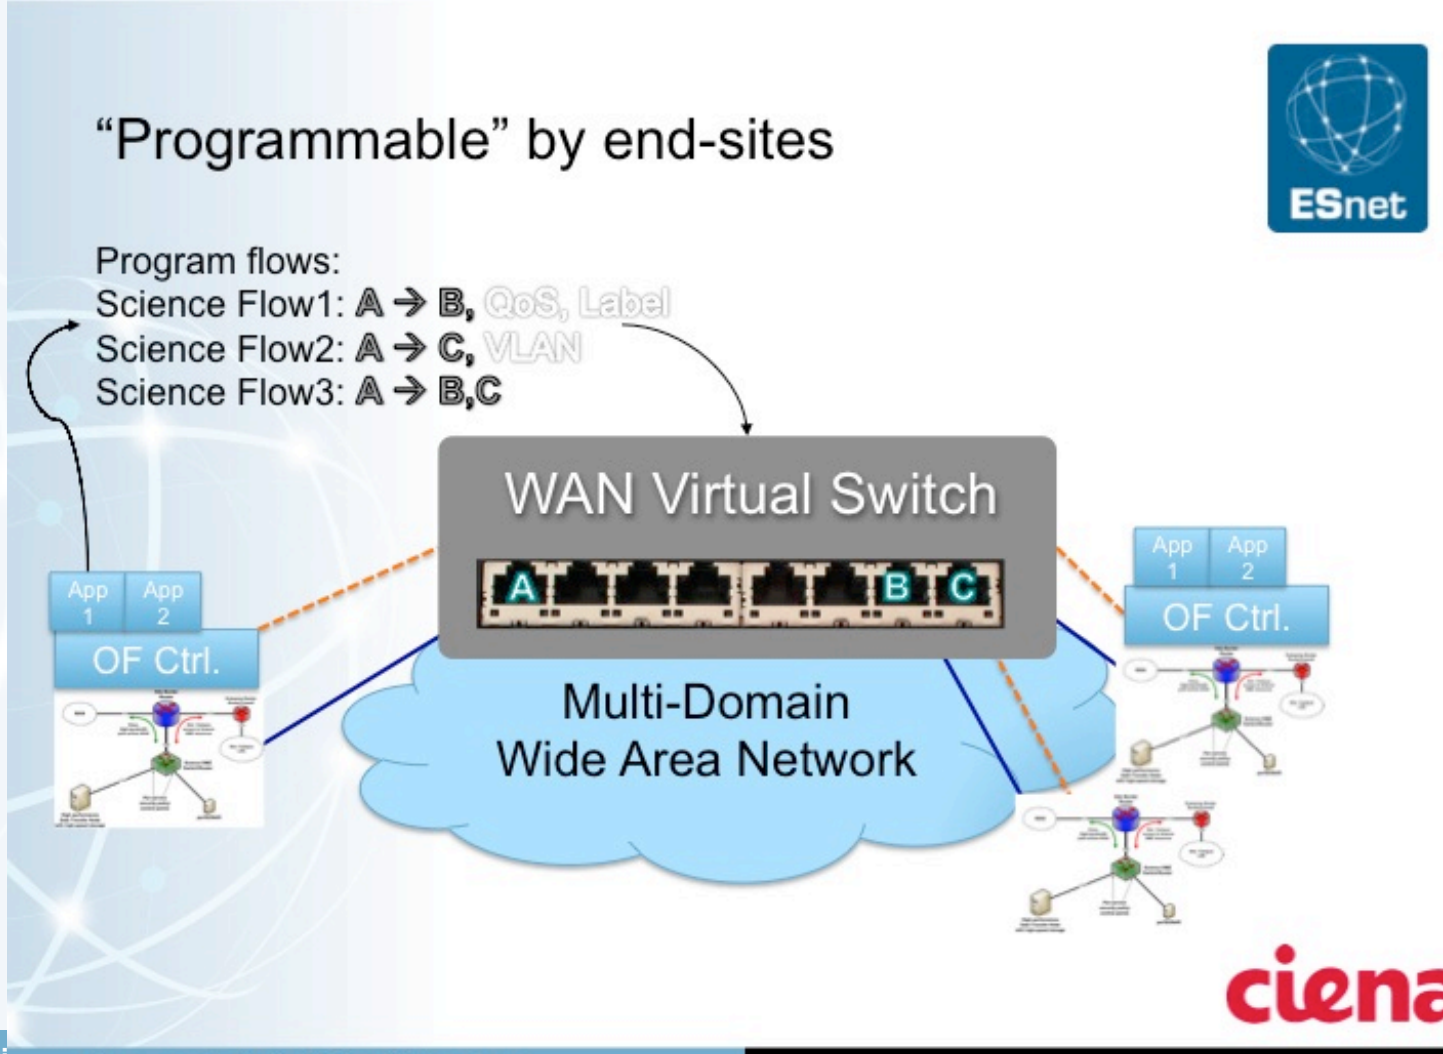

Journey with Programmability

World's first Transport SDN Demo, Infinera/ESnet/Brookhaven
December 2012

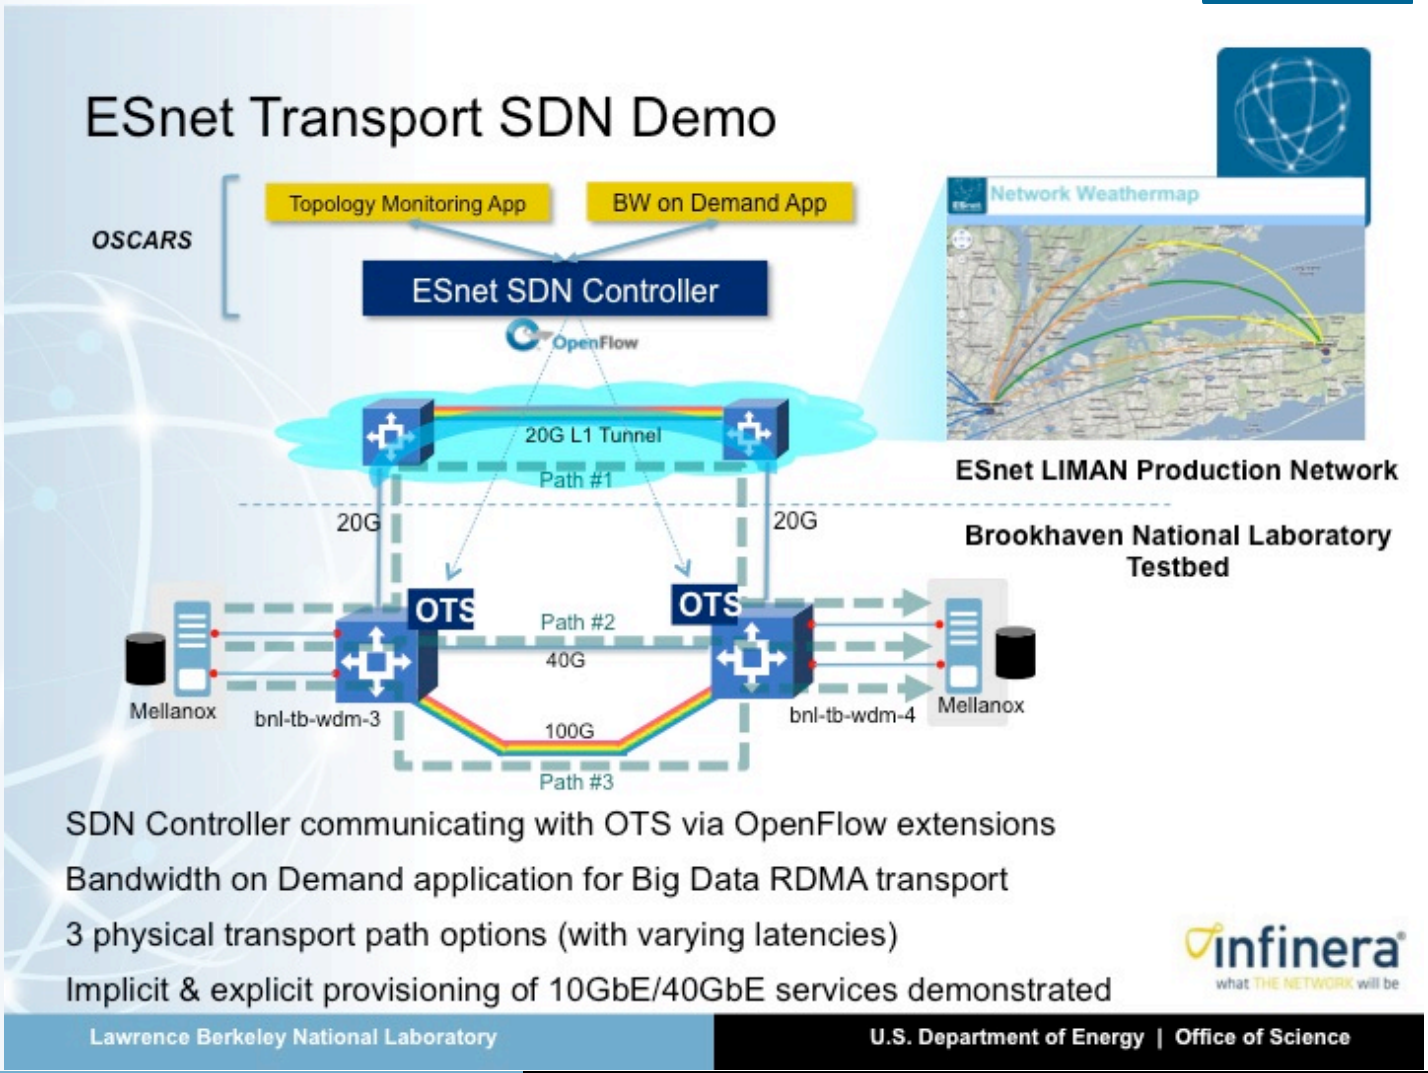

Insights

- optical-layer automation essential for future topologies, architectures

Industry impact – led to the formation of Optical Transport Working Group (OTWG) in ONF

- Feasibility of using OpenFlow for controlling optical devices justified formalization of the OTWG in ONF
- Invited to be a ONF Research Associate
  - ONLY DOE/Lab person represented in the ONF (with 200 companies)
- Multiple companies working on implementing the changes – big demonstration in Germany in October timeframe
  - ESnet not participating in demonstration due to resource constraints

Journey with Programmability
BGP over SDN infrastructure, ONS 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

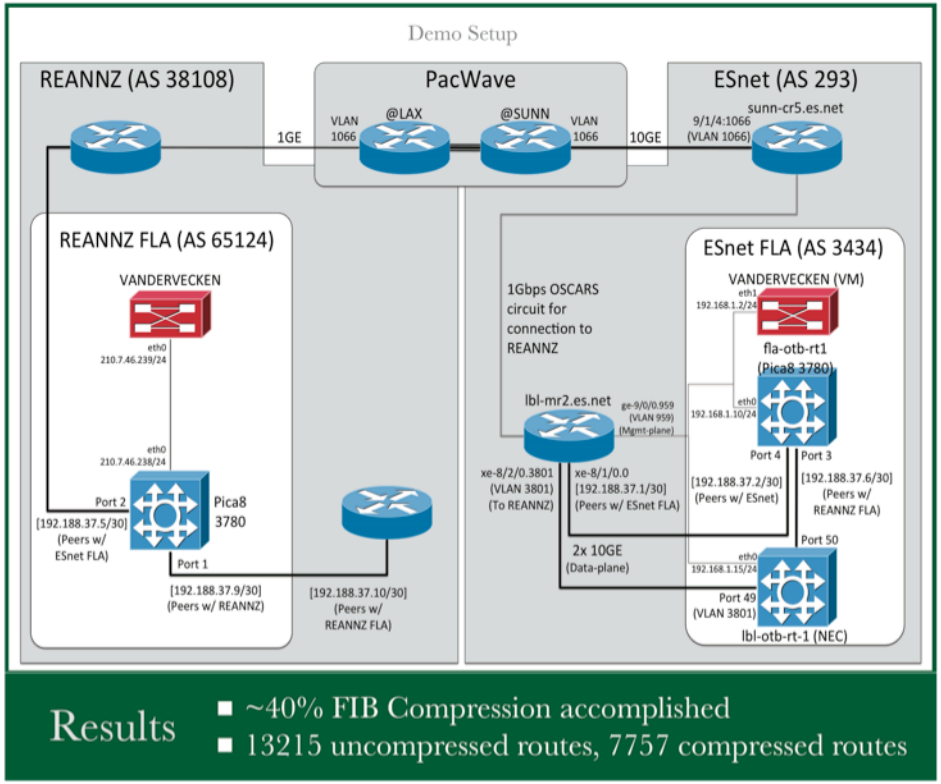

Insights

- SDN networks can interface with existing Internet
- New techniques need to be developed to scale controller based networking

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

Software-Defined Exchanges

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

A lot of hype on this exchange concept

Allow multiple networks to build an exchange and apply application/ network specific policy on traffic being exchanged

Definition being adapted by other folks to suit their purposes.

Journey with Programmability
Multi-layer SDN, Layer123 SDN, Oct 2013

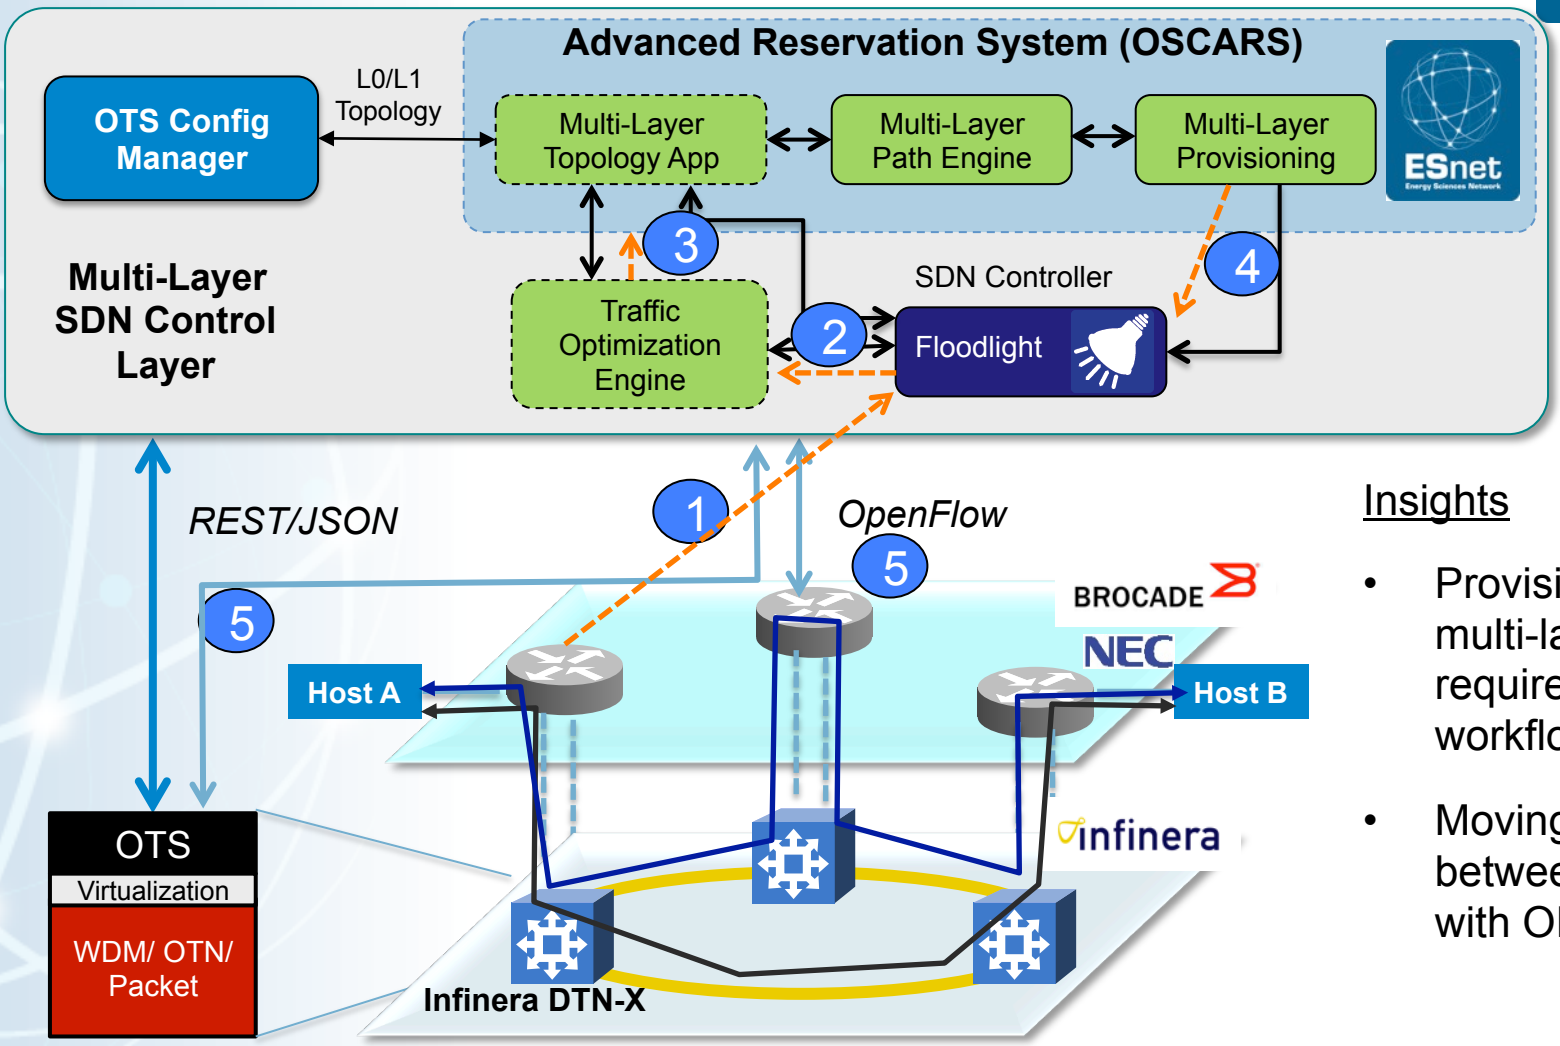

Insights

- Provisioning multi-layer could require a workflow
- Moving traffic between tunnels with OF is trivial

SDN Testbed Final Plan
- *not all the components are installed*

ESnet

ESnet OpenFlow Testbed

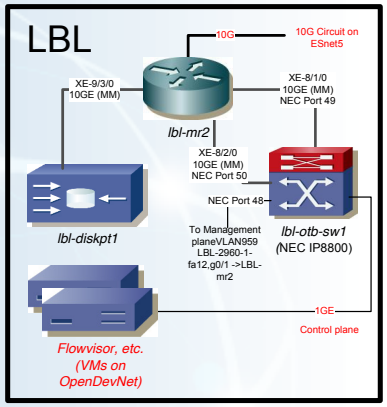

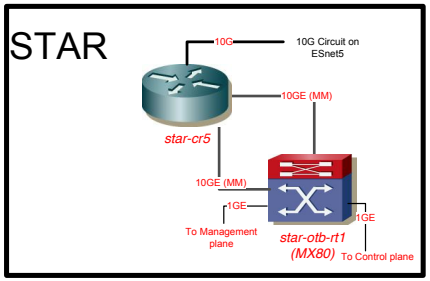

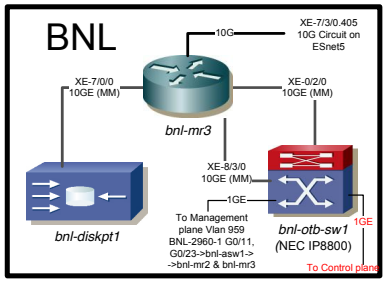

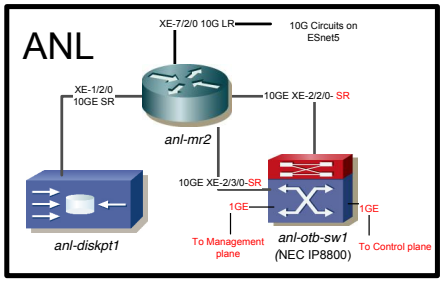

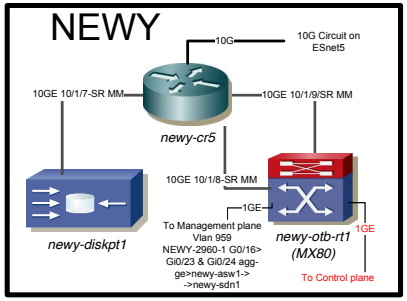

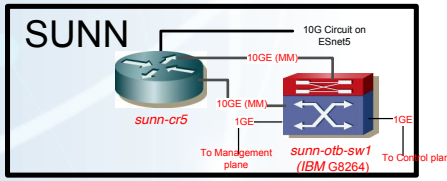

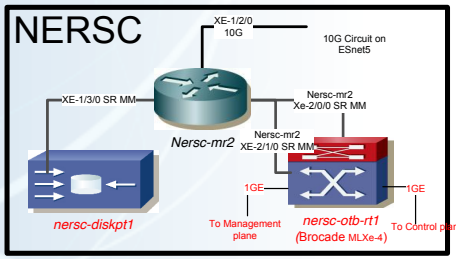

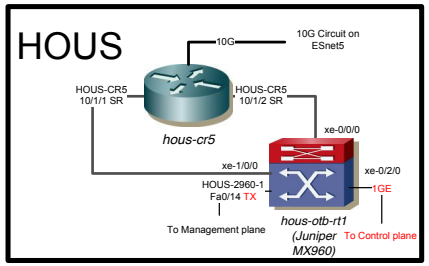

Updated February 26, 2014

OSCARS and SDN

Journey with Programmability

World's first Transport SDN Demo, Infinera/ESnet/Brookhaven

- OSCARS integrated OpenFlow as a PSS driver
- Vendor extensions to be compatible with OTS

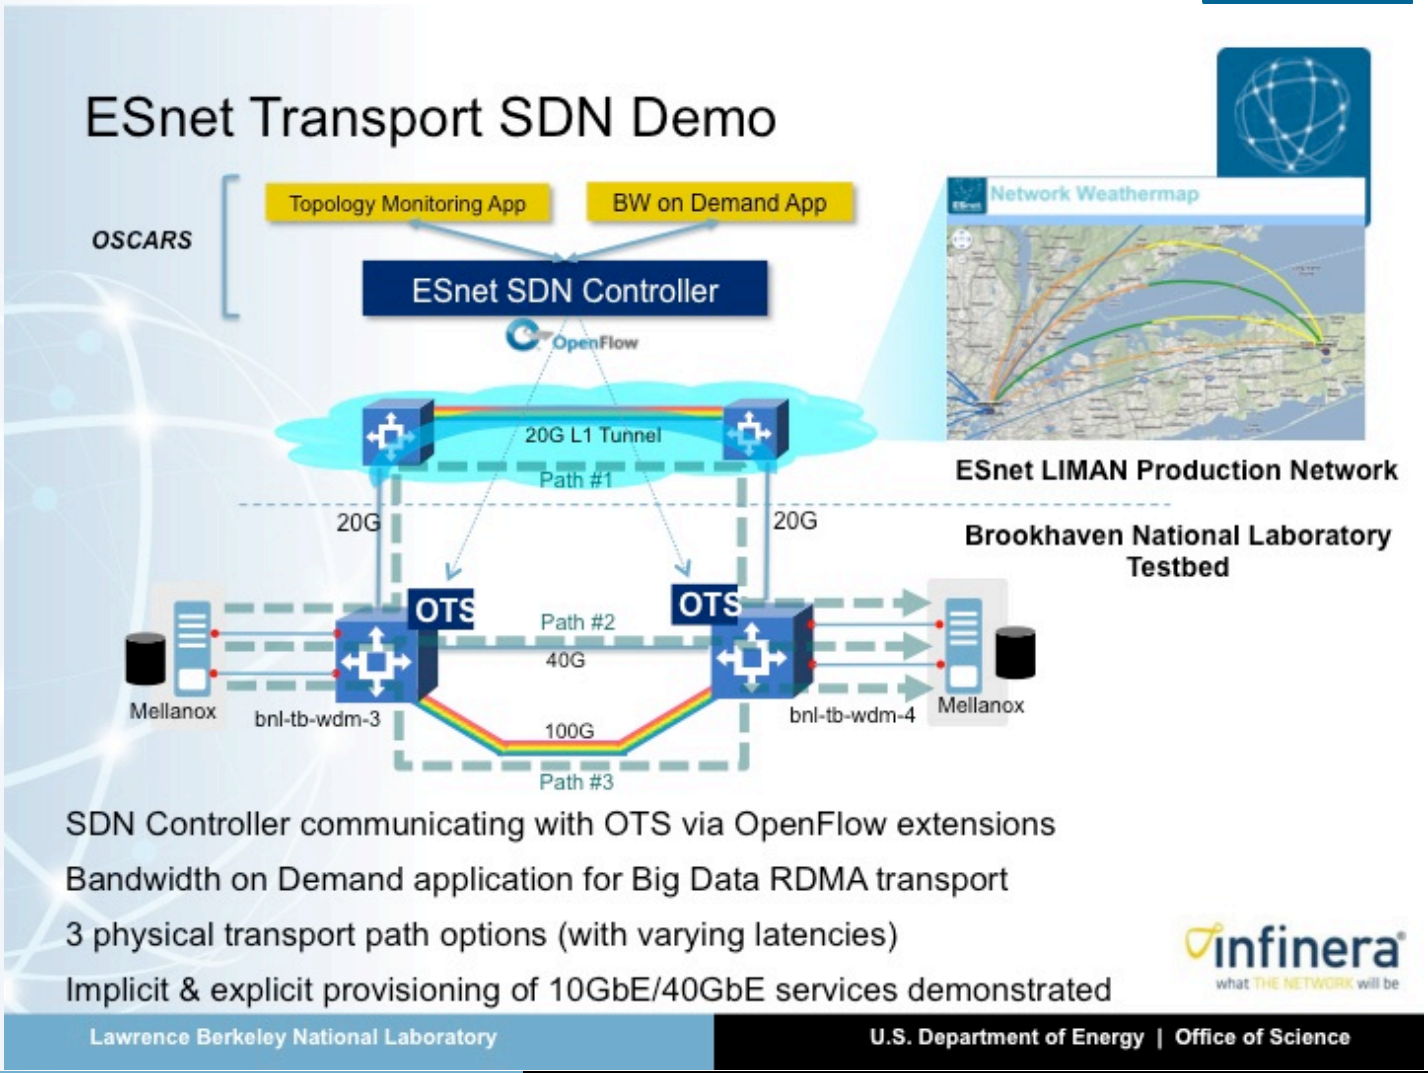

Journey with Programmability
Multi-layer SDN, Layer123 SDN, Oct 2013

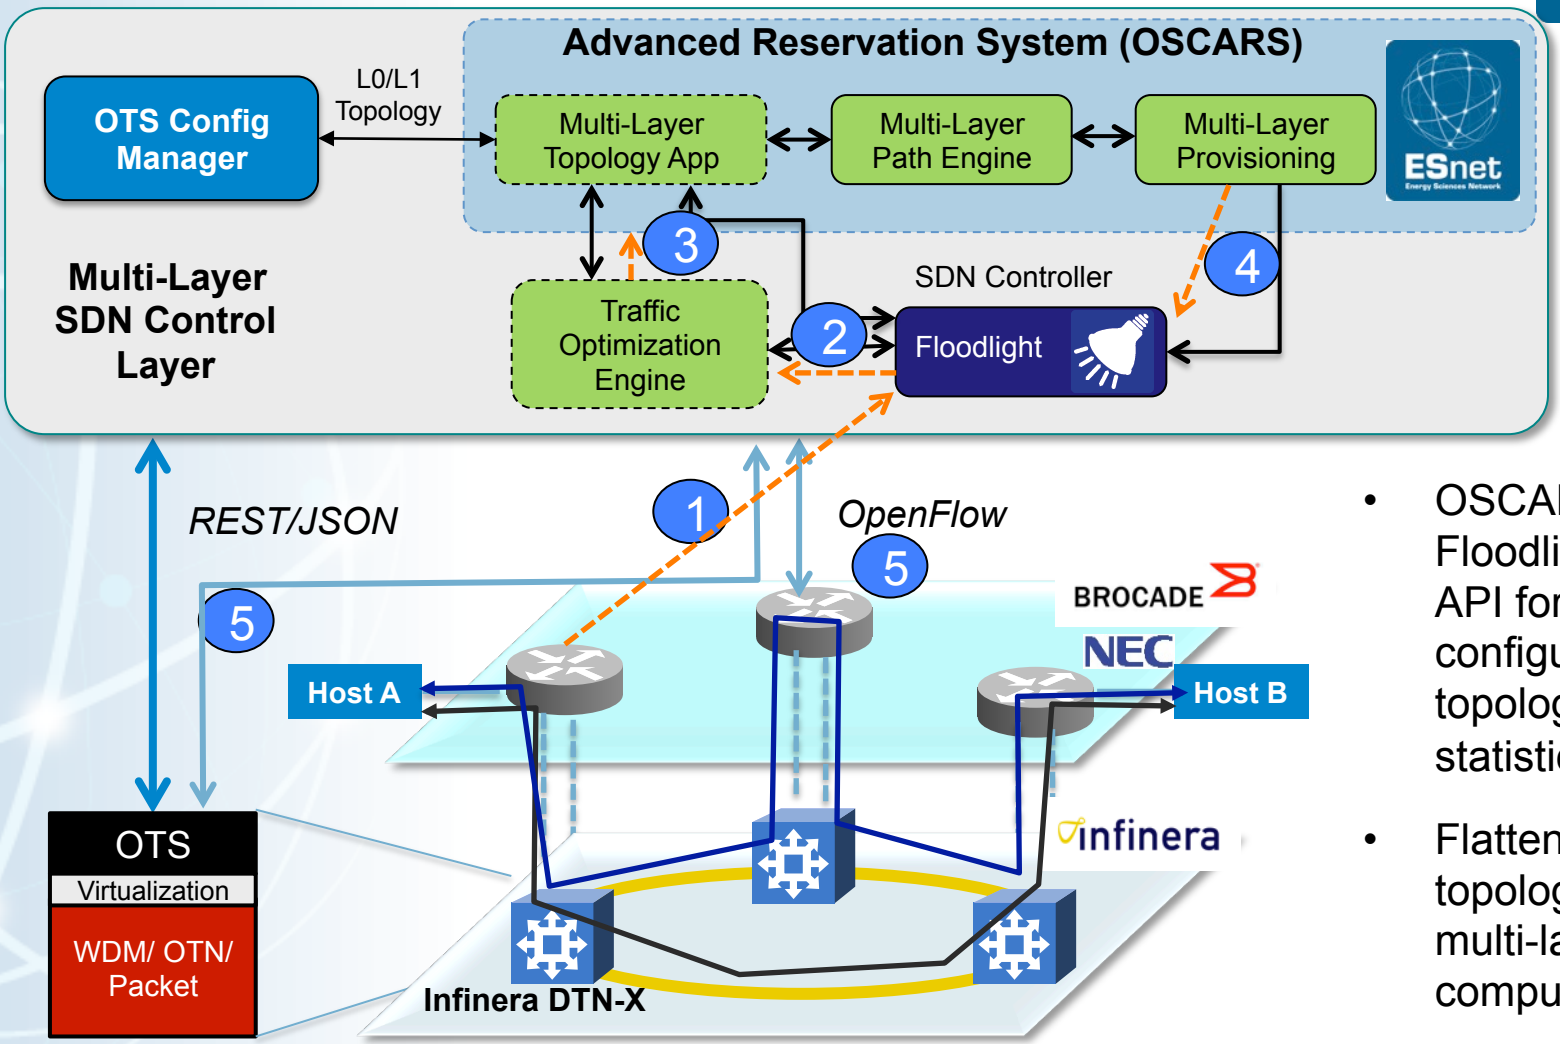

- OSCARS used Floodlight REST API for configuration, topology and statistics
- Flatten the topology for multi-layer path computation

OGF/ESNet - Service Request NBI

Allow simple service request across domain

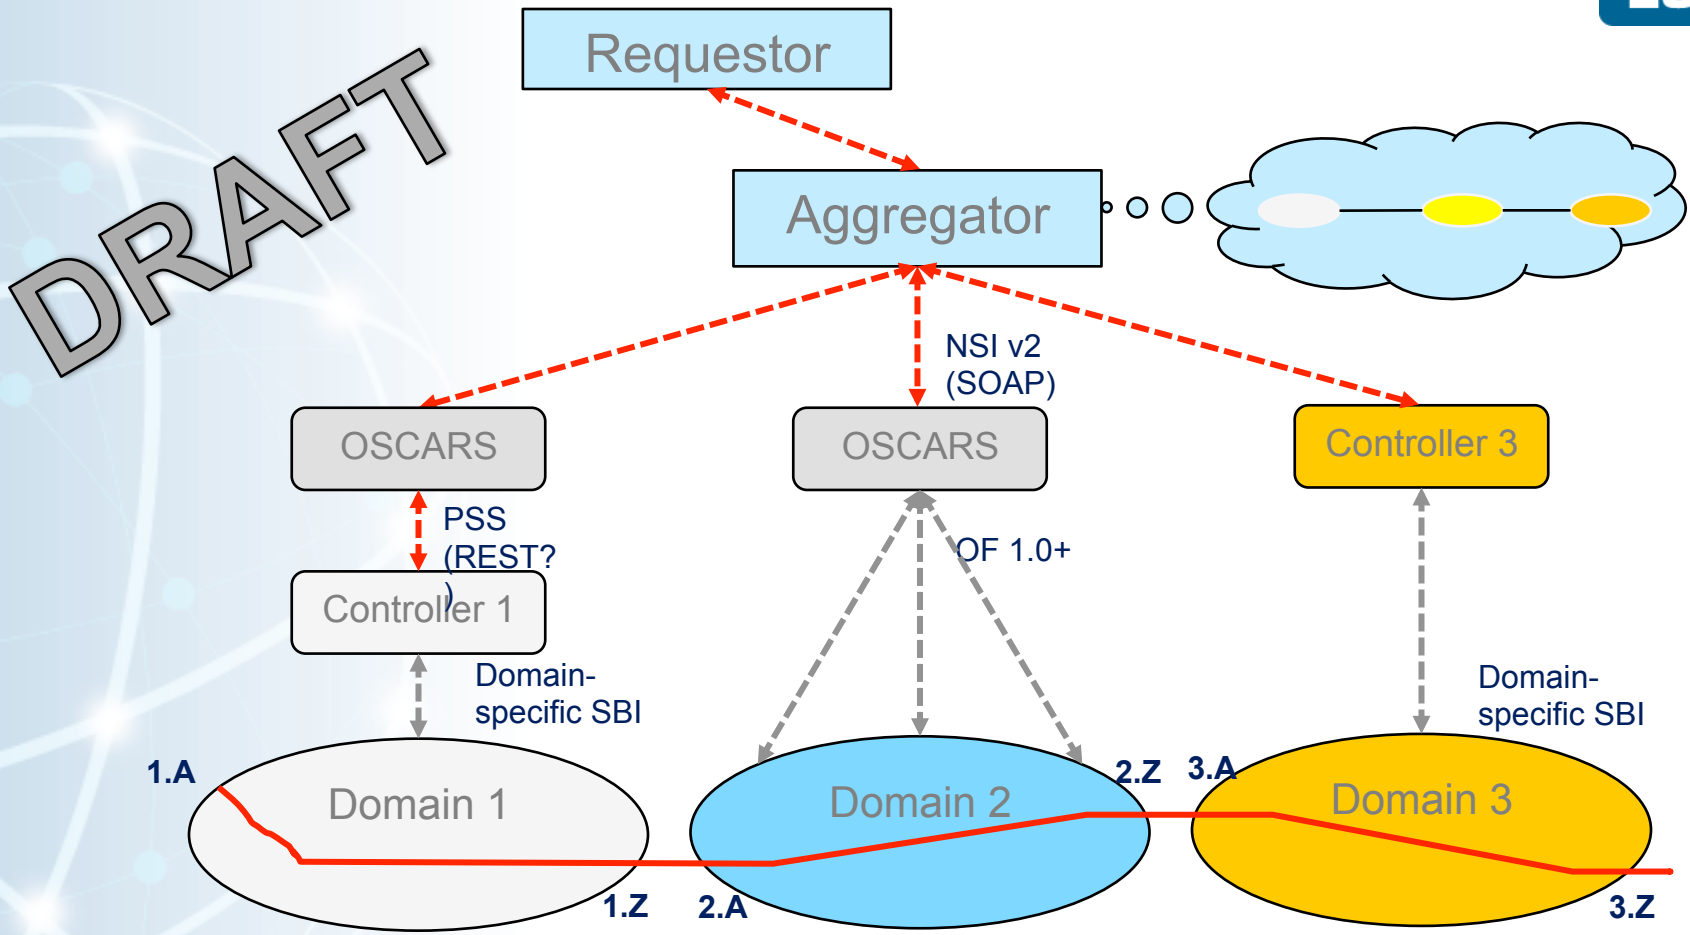

Building Network Capabilities using Atomic and Composite Network Services

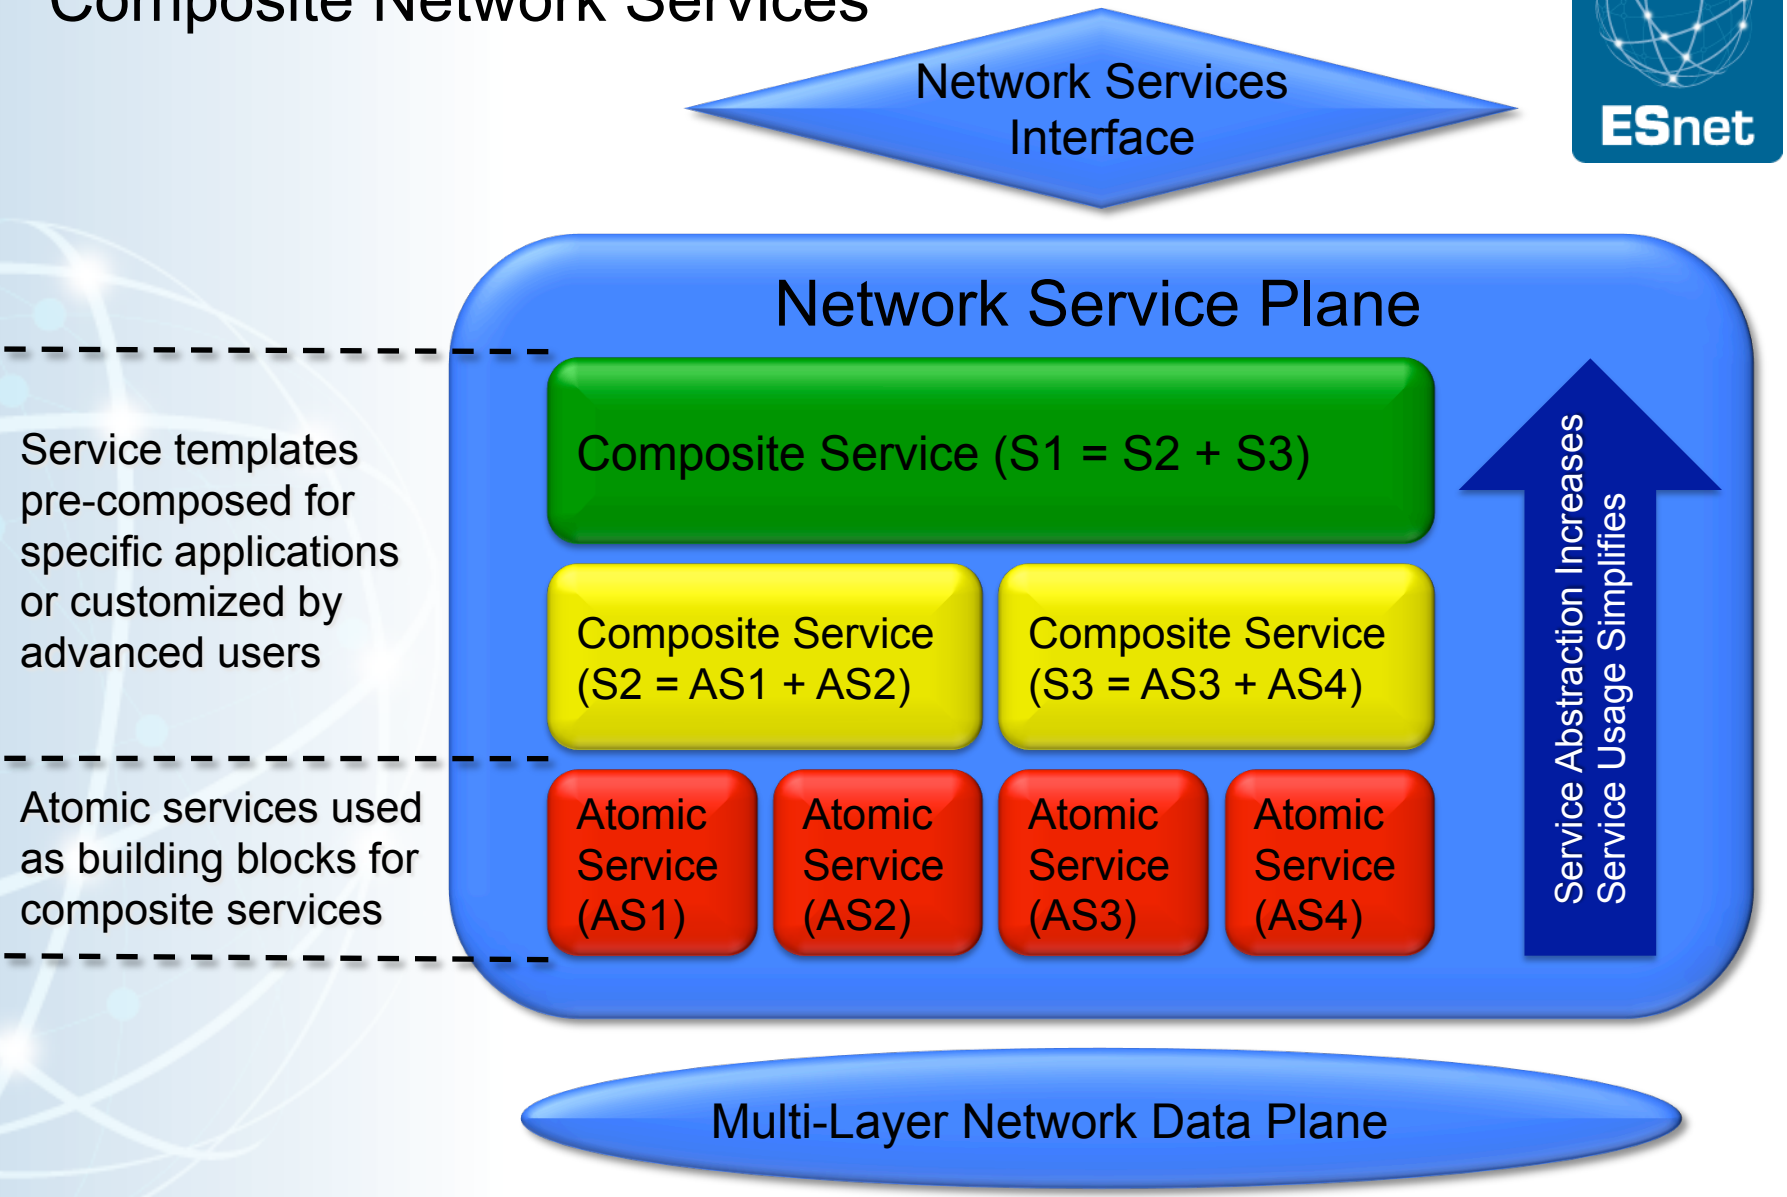

Examples of Atomic Services

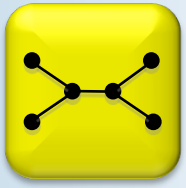

Topology Service to determine resources and orientation

Routing Service to enable IP connectivity

Resource Computation Service to determine possible resources based on multi-dimensional constraints

Store and Forward Service to enable caching capability in the network

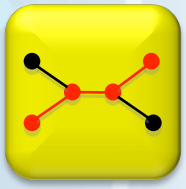

Connection Service to specify data plane connectivity

Measurement Service to enable collection of usage data and performance stats

Protection Service to enable resiliency through redundancy

Monitoring Service to ensure proper support using SOPs for production service

Restoration Service to facilitate recovery

Firewall Service to prevent unwanted access of network resources

Examples of Composite Network Services

LHC: Resilient High Bandwidth Guaranteed Connection


Connection

Protection


Measurement


Monitoring

L3VPN: Privately Routed Network (w/ Commodity)


Connection


Restoration


Routing


Firewall


Monitorin

Protocol Testing: Constrained Path Connection


Topology


Resource Computation


Connection


Measurement

Conclusion: Application Workflow Integration is Critical!

A key focus is on technology development which allow networks to participate in application workflows

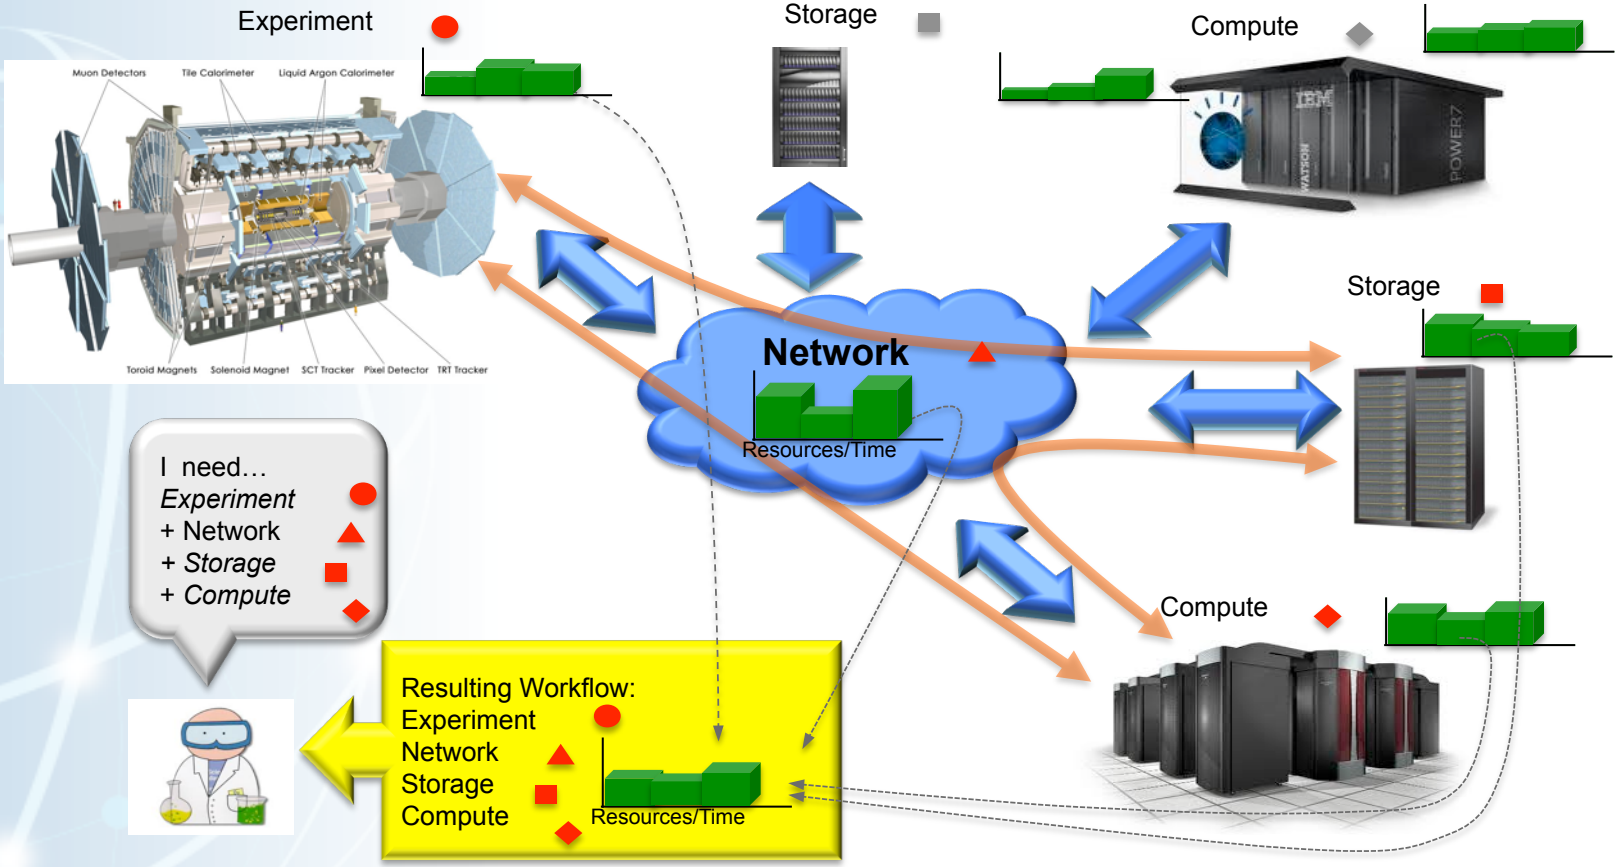

The Network needs to be available to application workflows as a first class resource in this ecosystem

DOE/NSF Workshop
‘SDN Operationalization’

Three Goals

1. Bridge the ‘operational’ gap
   - architecture, tools and policies
2. Deploy and operate **securely** multi-layer, multi-domain SDN networks
   - Interwork with the current set of Internet technologies
3. Identify research, development and technologies needed to support new, innovative users and applications

It is about *running* networks (not just about network *research*)

Develop, deploy and (inter)operate
a prototype multi-domain SDN network

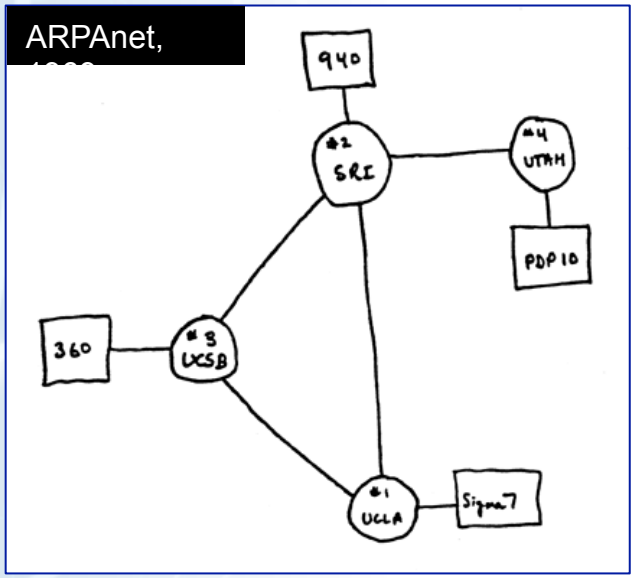

http://www.computerhistory.org/internet_history/full_size_images/1969_4-node_map.gif

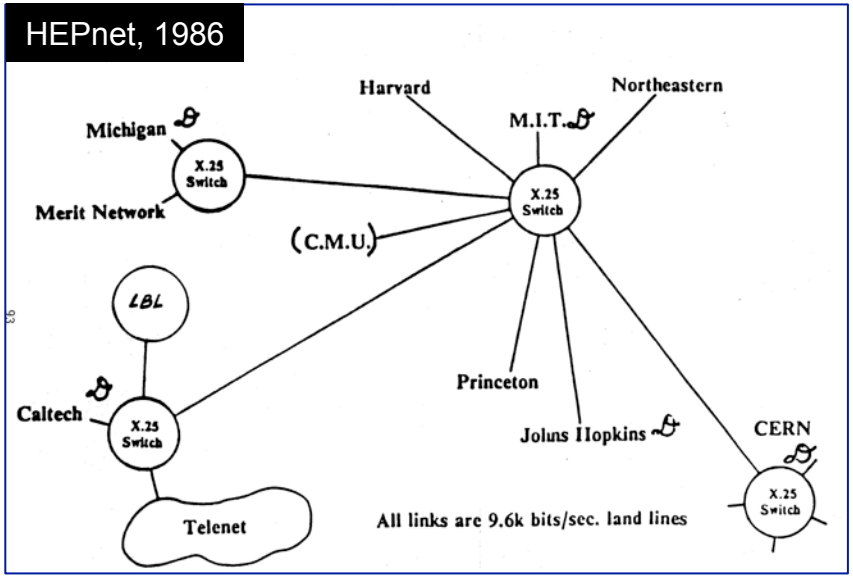

Future of intersite networking, LBL, 1986

Workshop to focus on the gaps and explore building a prototype multi-domain SDN network

1. Bridge the ‘operational’ gap
   - architecture, tools and policies
2. Deploy and operate **securely** multi-layer, multi-domain SDN networks
   - Interwork with the current set of Internet technologies
3. Identify research, development and technologies needed to support new, innovative users and applications

Observations and gaps (only a subset)

Time is right for prototyping operational, multi-domain SDNs

- Connect up testbeds and operational networks
- Build the SBone, the ARPAnet, ….with new technologies

Multi-domain SDN deployment is key

Observations and gaps contd.

Build multiple, community supported Software-Defined eXchanges (SDX)

- Tackling exchanging IP data is table-stakes for a larger deployment

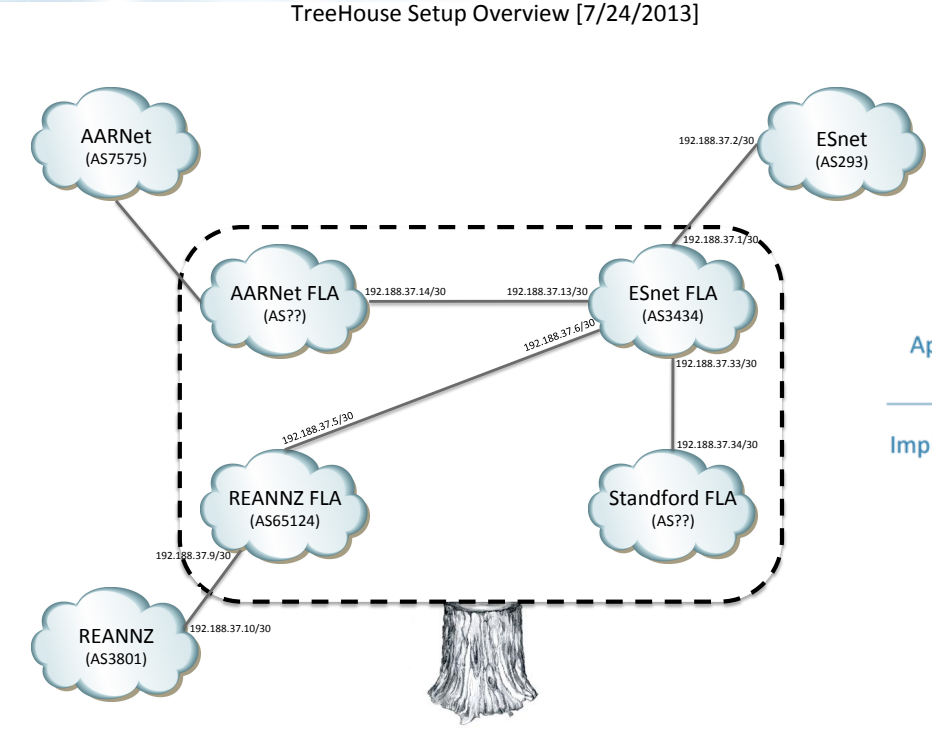

Collaboration with Josh Bailey, Google

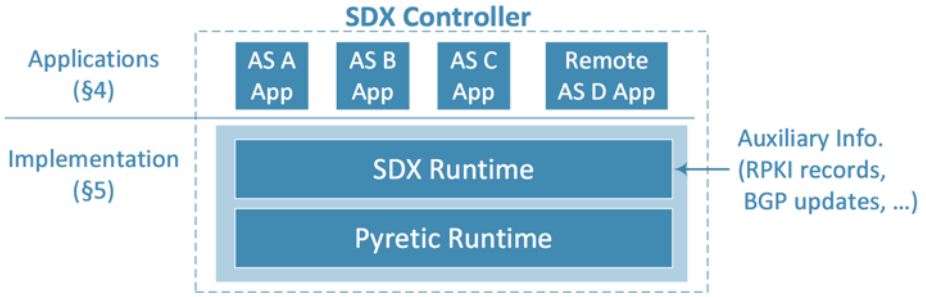

Nick Feamster, Georgial Tech