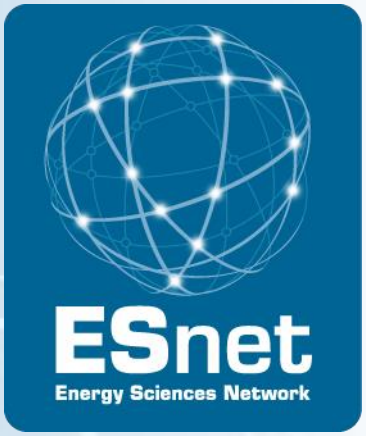

Software-Defined Networks (SDN): Bridging the application-network divide

Inder Monga

Chief Technologist and Area Lead,

Engineering, Research and Software development

CHEP 2013

Shared historical moment

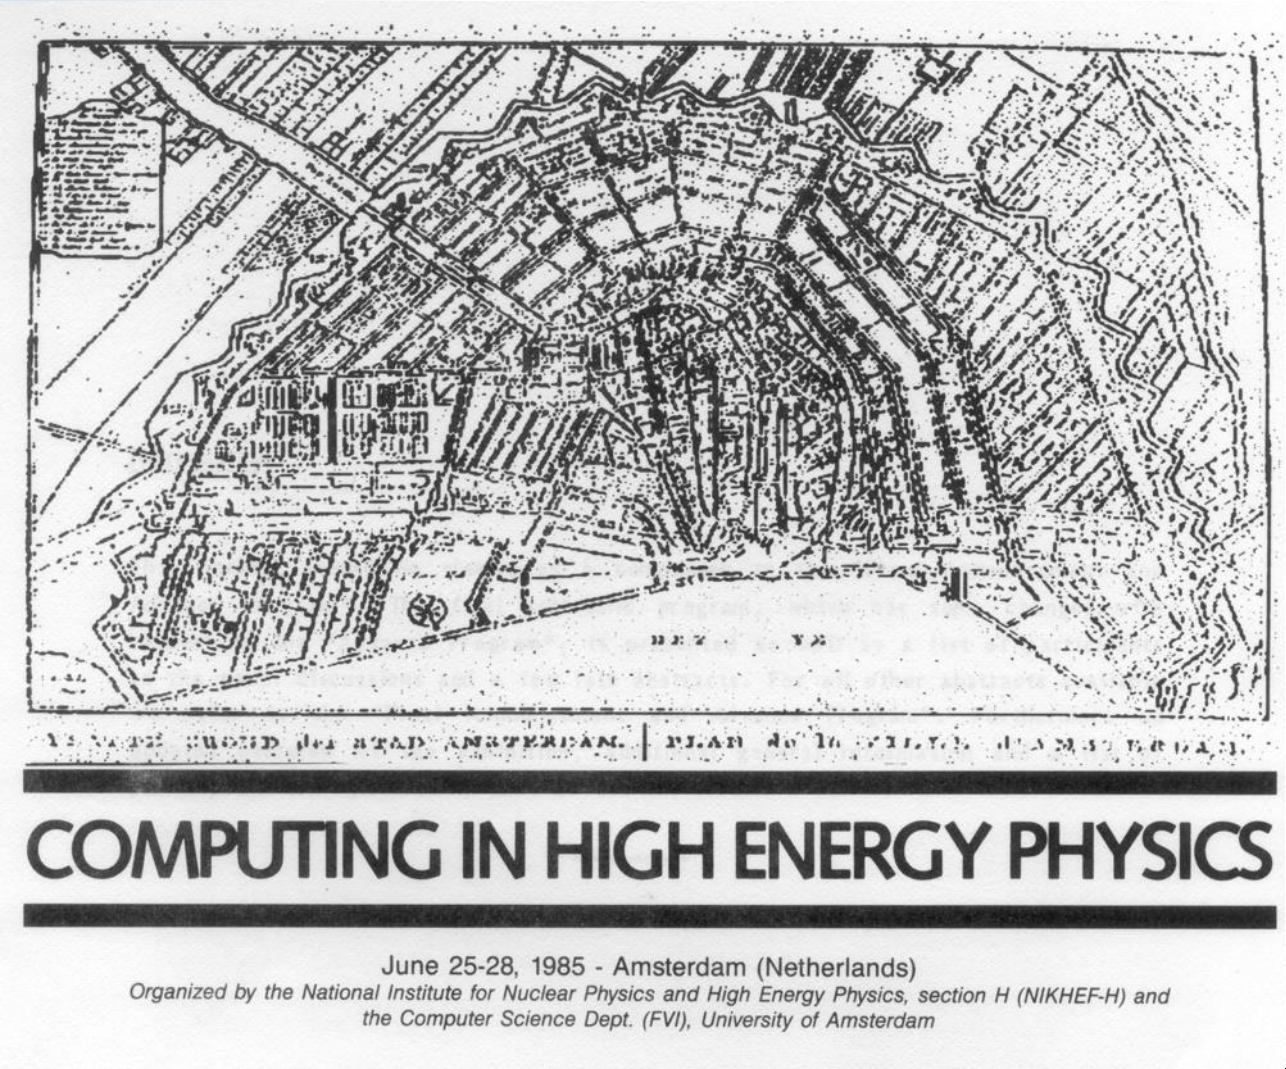

COMPUTING IN HIGH ENERGY PHYSICS

June 25-28, 1985 - Amsterdam (Netherlands)

Organized by the National Institute for Nuclear Physics and High Energy Physics, section H (NIKHEF-H) and the Computer Science Dept. (FVI), University of Amsterdam

Thank you! HEP community helped create ESnet 27 years ago

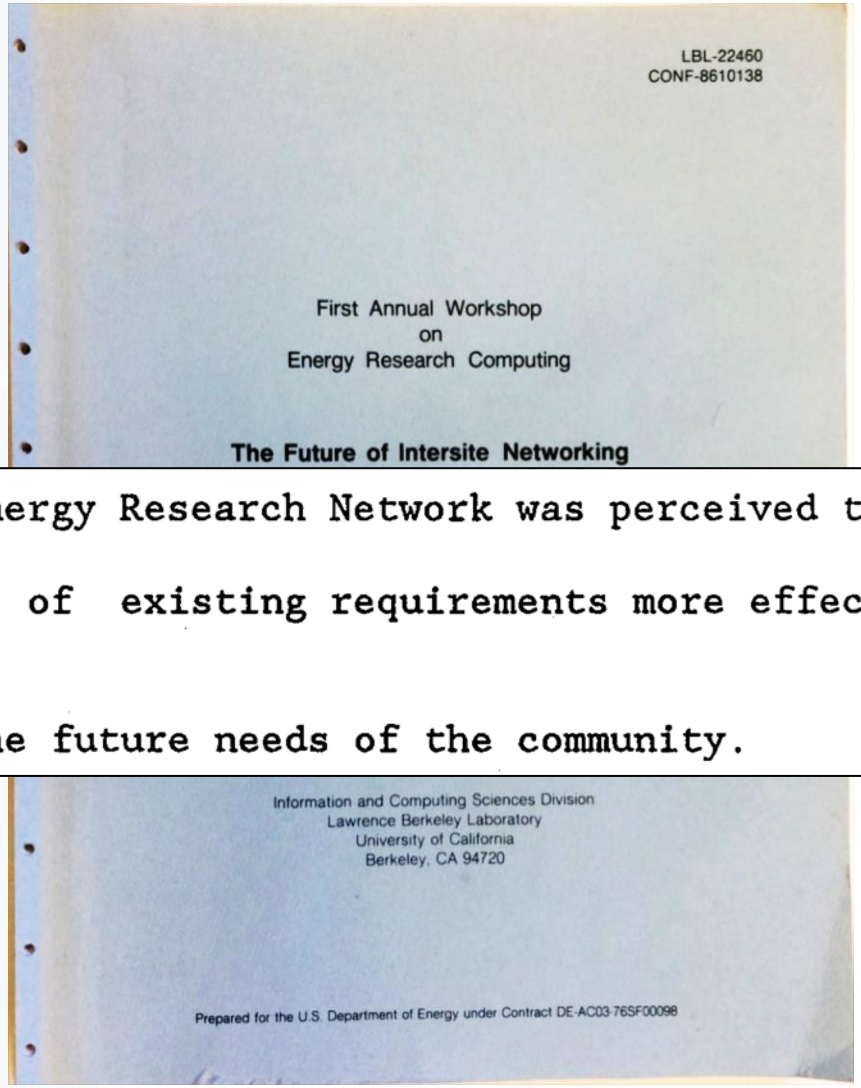

LBL-22460
CONF-8610138

First Annual Workshop
on
Energy Research Computing

The Future of Intersite Networking

Information and Computing Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

Prepared for the U.S. Department of Energy under Contract DE-AC03-76SF00098

The task of the Energy Research Network was perceived to be:

1) The serving of existing requirements more effectively and for less cost.

2) To look to the future needs of the community.

Science Data Transferred Each Month by the Energy Sciences Network

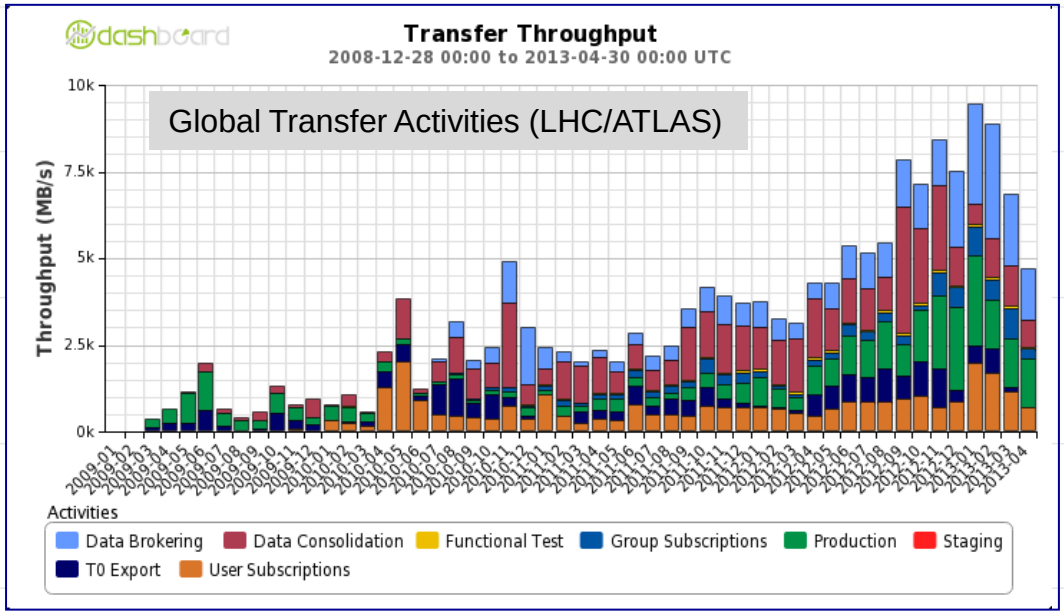

15.5 Petabytes
(March 2013)

Bytes Transferred

1.6×10^{16}
1.4×10^{16}
1.2×10^{16}
1×10^{16}
8×10^{15}
6×10^{15}
4×10^{15}
2×10^{15}
0

Jan-90
Feb-94
Apr-98
Jun-02
Aug-06
Oct-10

Month

27 years and 5 generations later….
World's first 100G Science Network

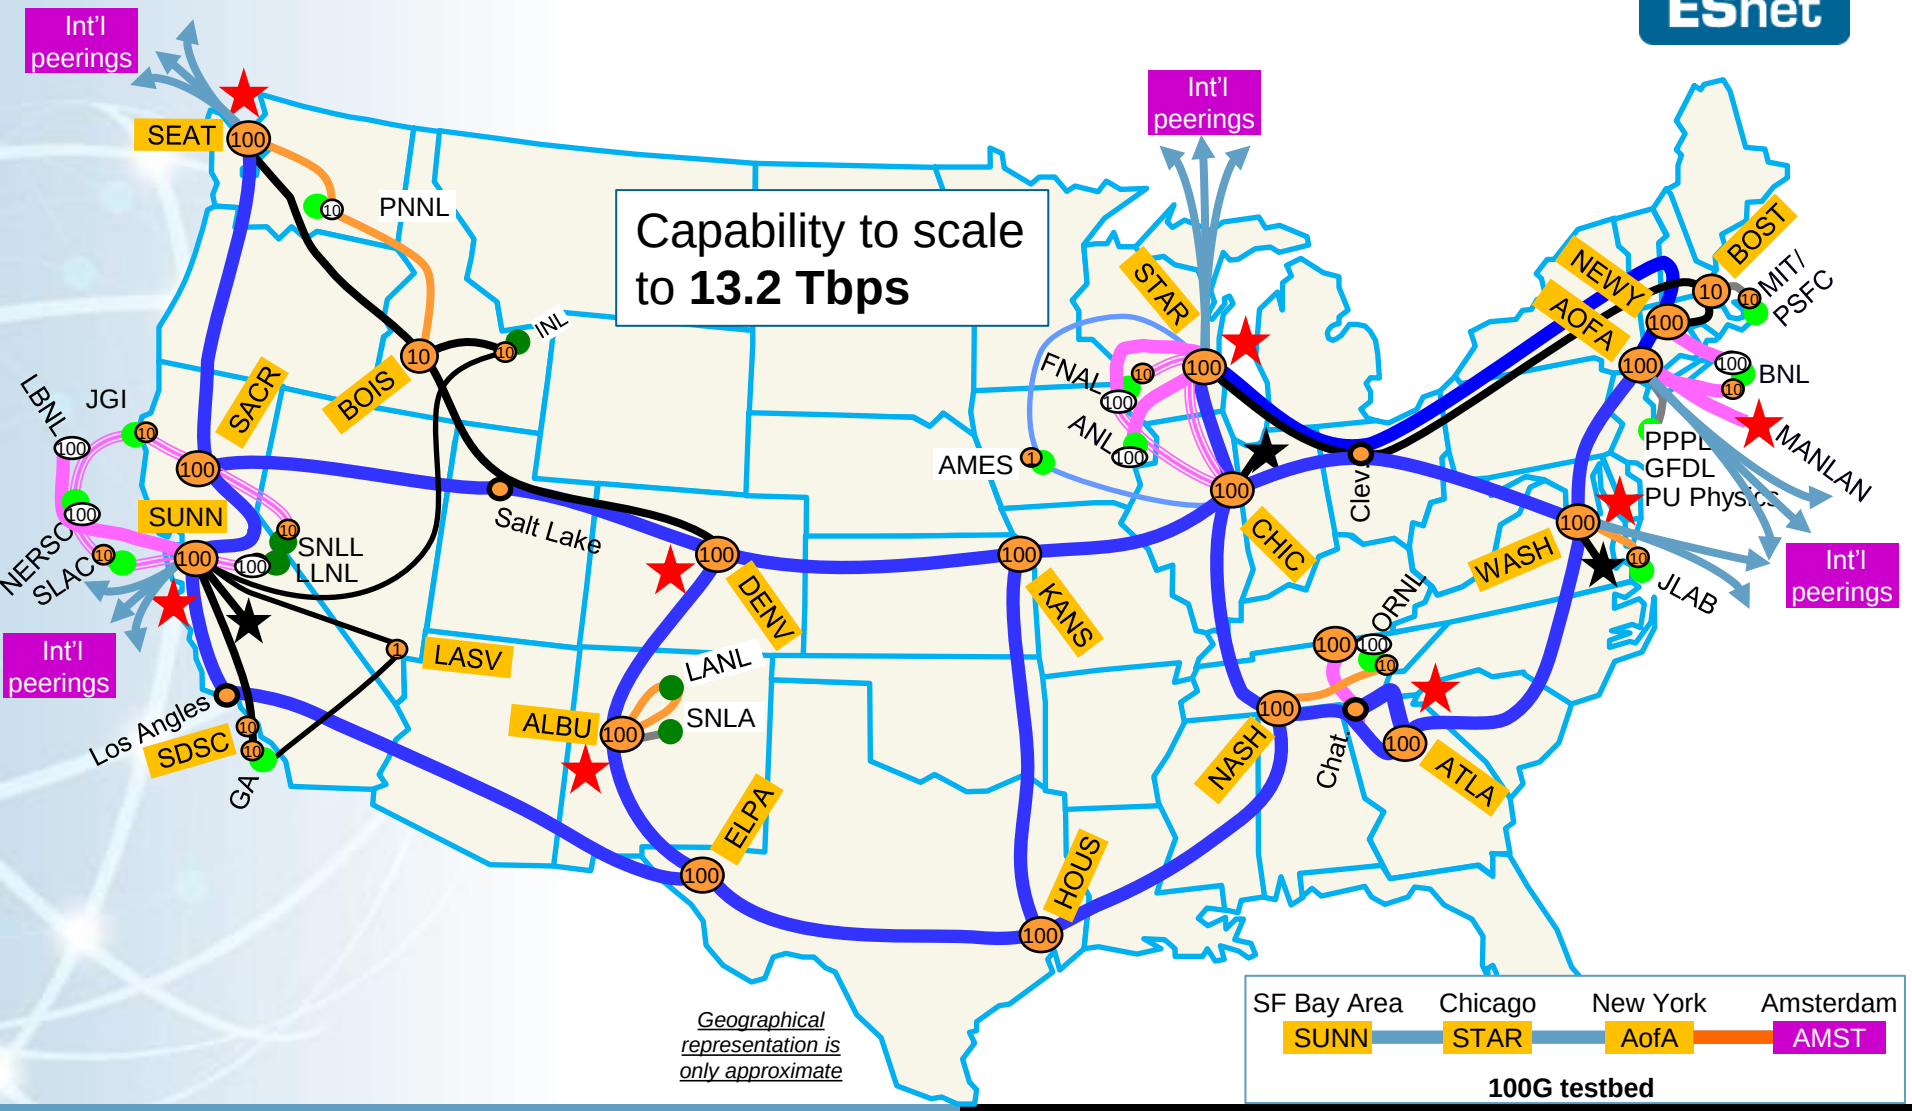

‘Back to the future’

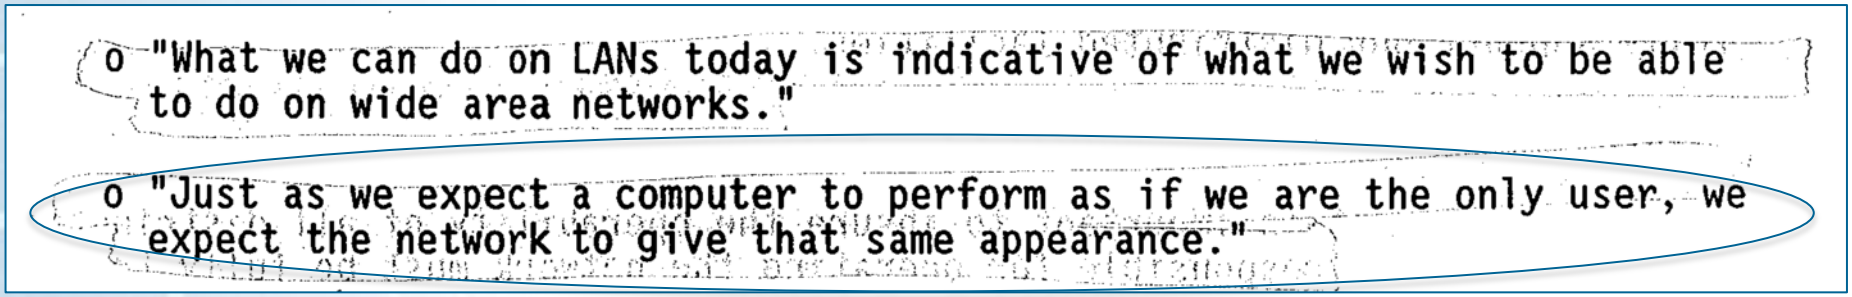

o "What we can do on LANs today is indicative of what we wish to be able to do on wide area networks."

o "Just as we expect a computer to perform as if we are the only user, we expect the network to give that same appearance."

First workshop report for ESnet on intersite networking, 1986

Network community is still struggling to meet application requirements captured in 1986!

Brute force approach (add more bandwidth) is not going to meet those requirements

What is common between modern networks and analog phone switches?

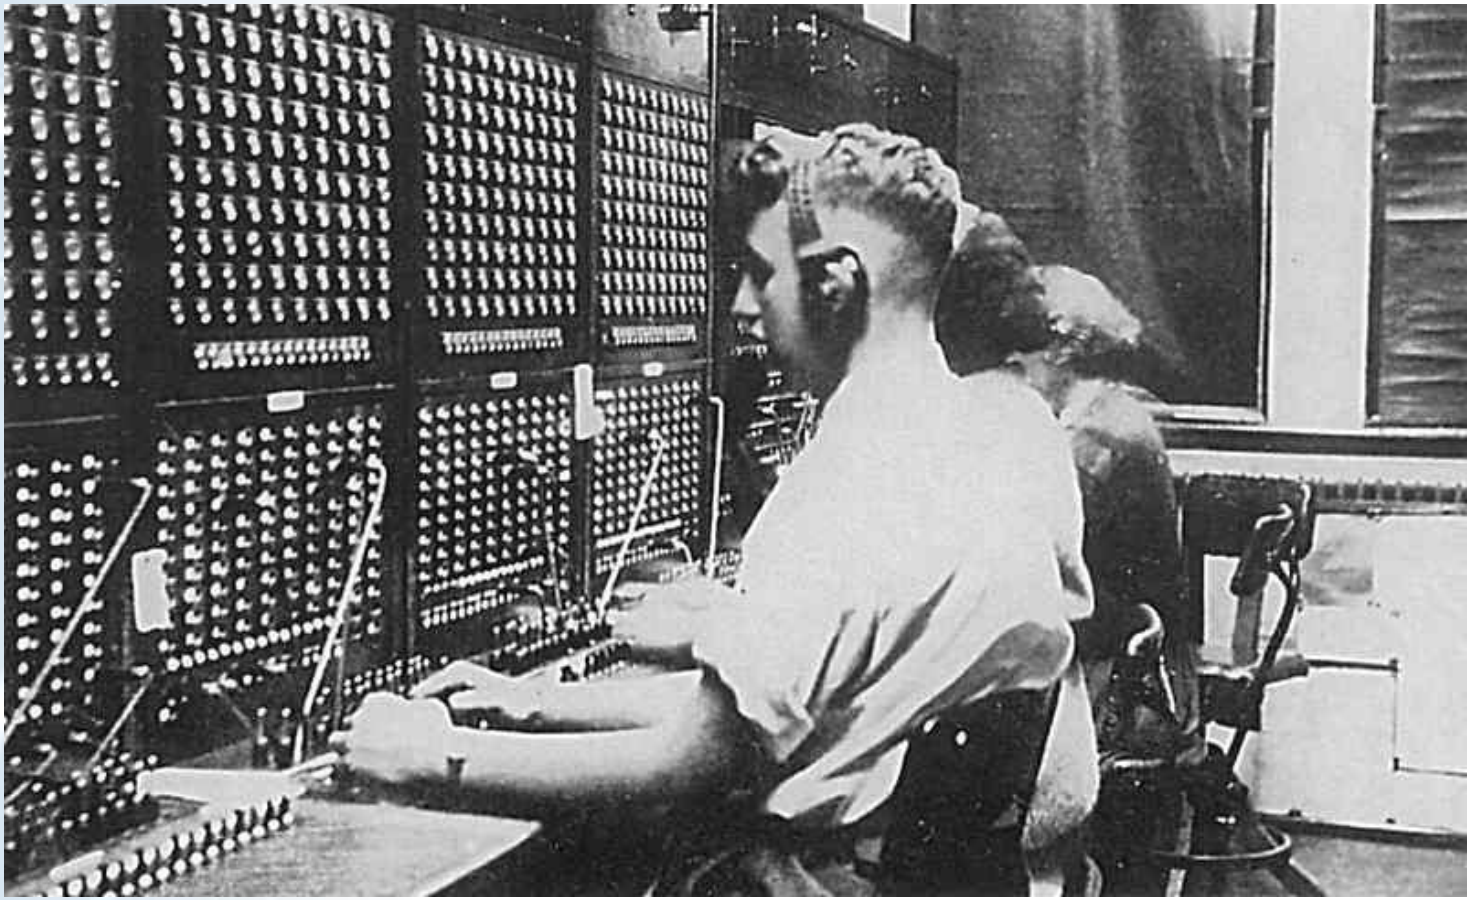

Labor-intensive, nearly static, error prone

```
Router(config)# class-map efcls
Router(config-cmap)# match ip dscp 46
Router(config-cmap)# exit
Router(config)# access-list 100 permit ip any any
Router(config)# class-map becls
Router(config-cmap)# match access-group 100
Router(config-cmap)# exit
Router(config)#
Router(config)# policy-map ef_prio
Router(config-pmap)# class efcls
Router(config-pmap-c)# priority 1500
Router(config-pmap-c)# exit
Router(config-pmap)# class becls
Router(config-pmap-c)# random-detect
Router(config-pmap-c)# random-detect precedence 0 15 45 3
Router(config-pmap-c)# random-detect precedence 1 15 45 3
Router(config-pmap-c)# random-detect precedence 2 15 45 3
Router(config-pmap-c)# random-detect precedence 3 15 45 3
Router(config-pmap-c)# random-detect precedence 4 15 45 3
Router(config-pmap-c)# random-detect precedence 5 15 45 3
Router(config-pmap-c)# random-detect precedence 6 15 45 3
Router(config-pmap-c)# random-detect precedence 7 15 45 3
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)#
```

cm-036

What is SDN?

Loose definition: separation of data-plane from control plane
In essence: enables programmability

programmable

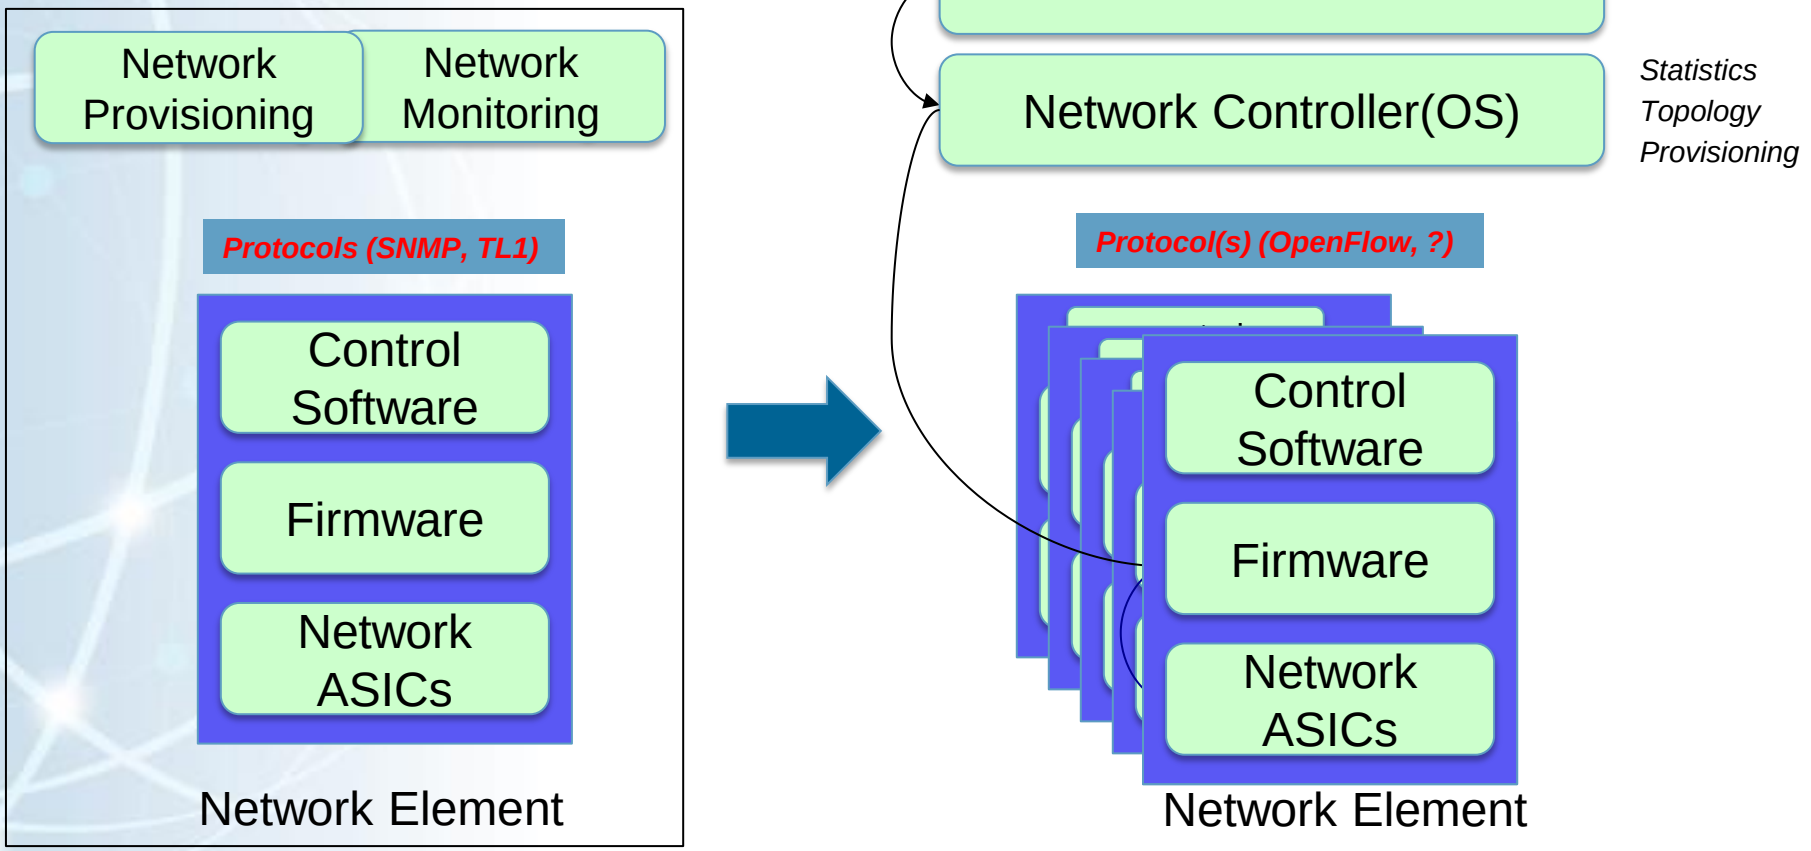

Simple programming constructs

OpenFlow 1.0 standard

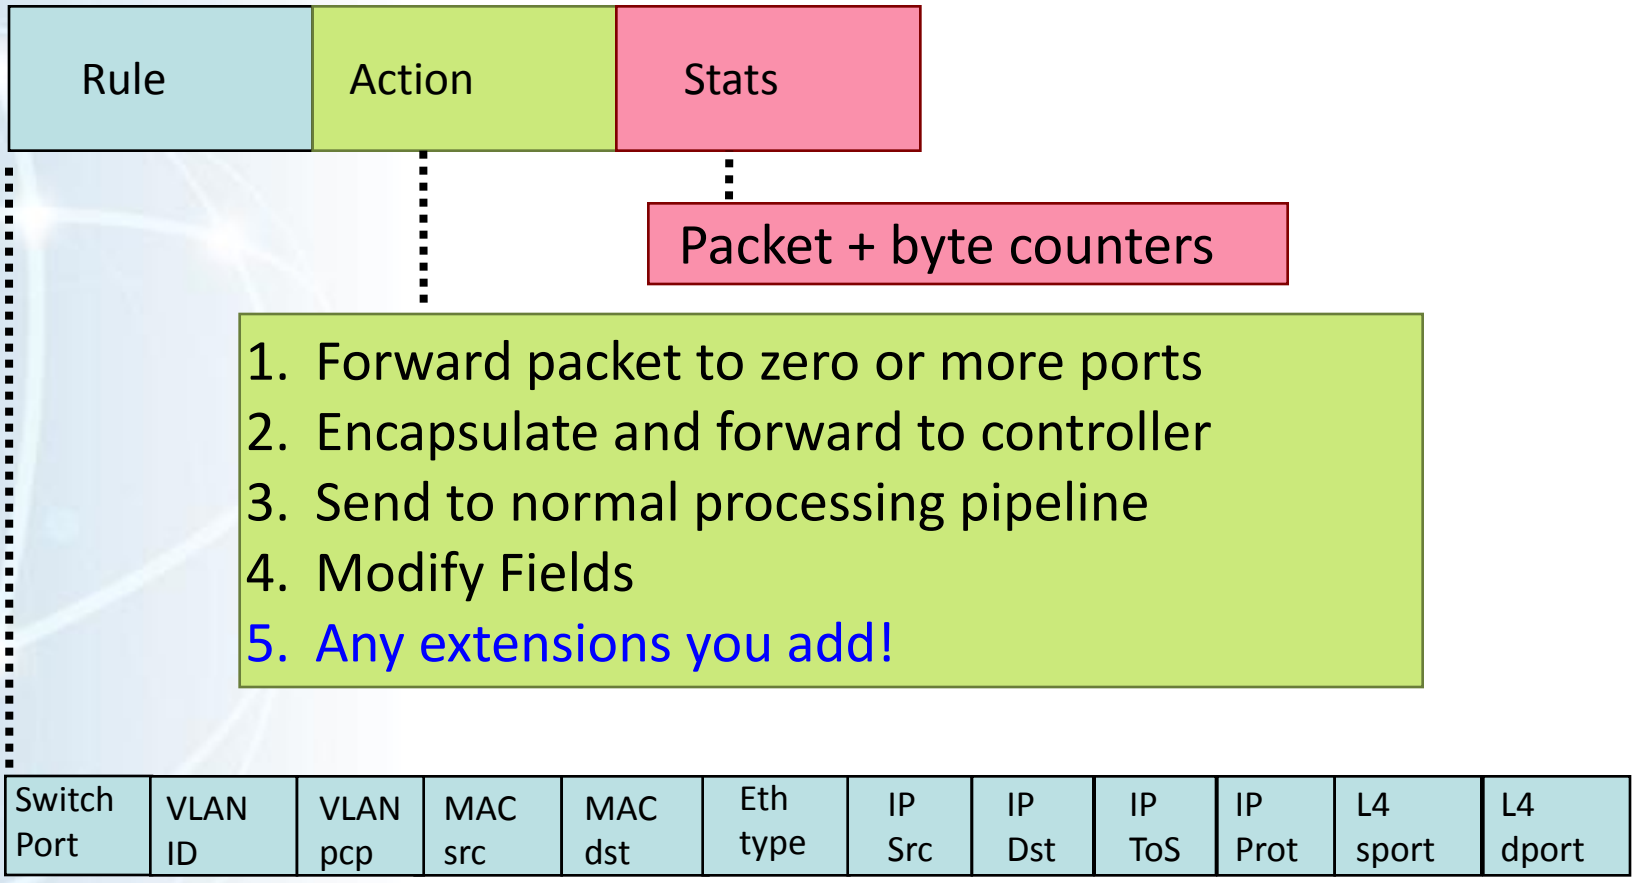

+ mask what fields to match

Slide courtesy Srini Seetharaman

OpenFlow usage
Classic model

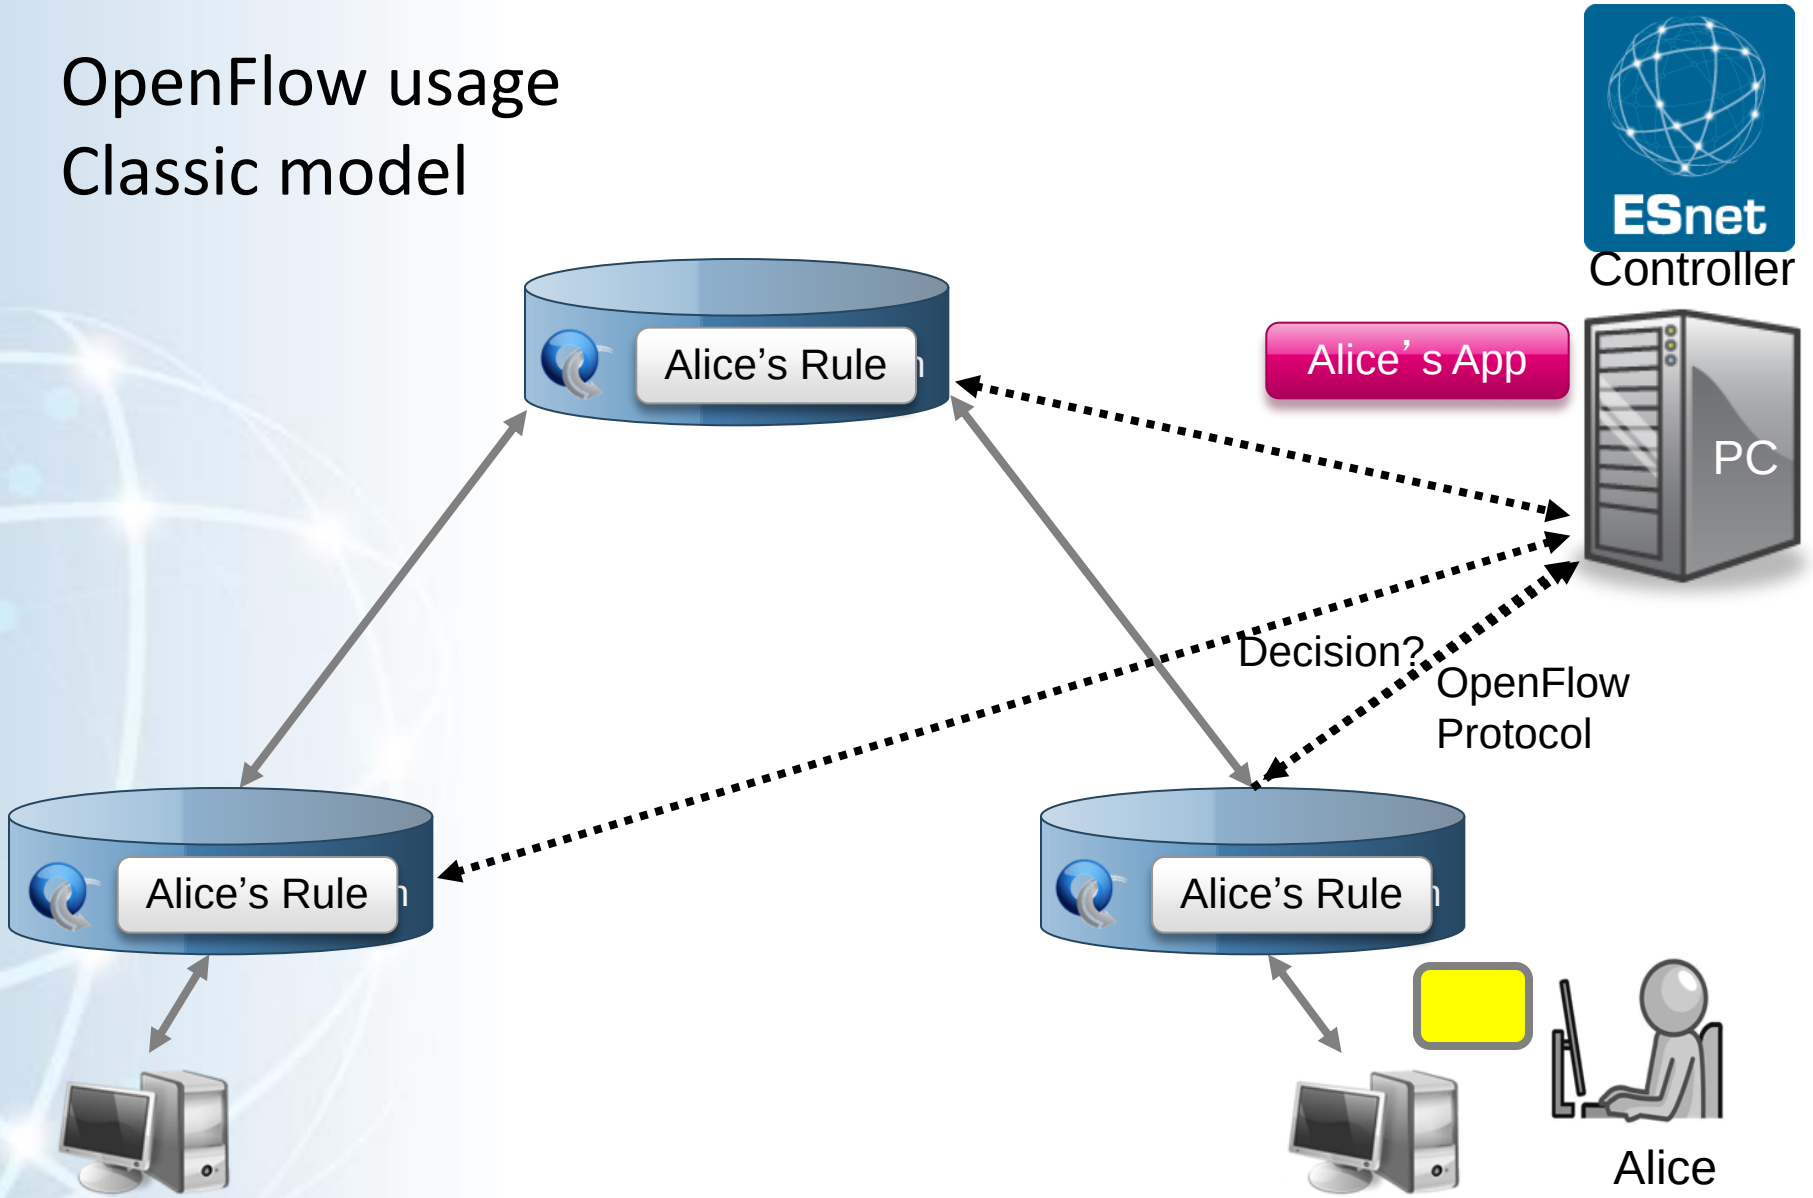

OpenFlow offloads control intelligence to remote software

What is the paradigm difference?

Internet today:
- Built-in control in each layer
- Multiple management domains

SDN Approach:
- Network-wide cross-layer view
- Unclear how management evolves

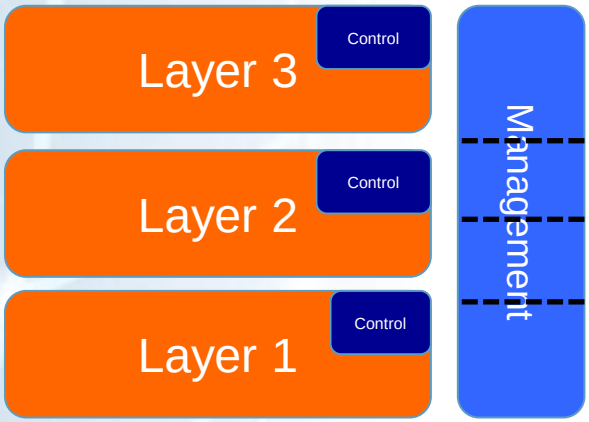

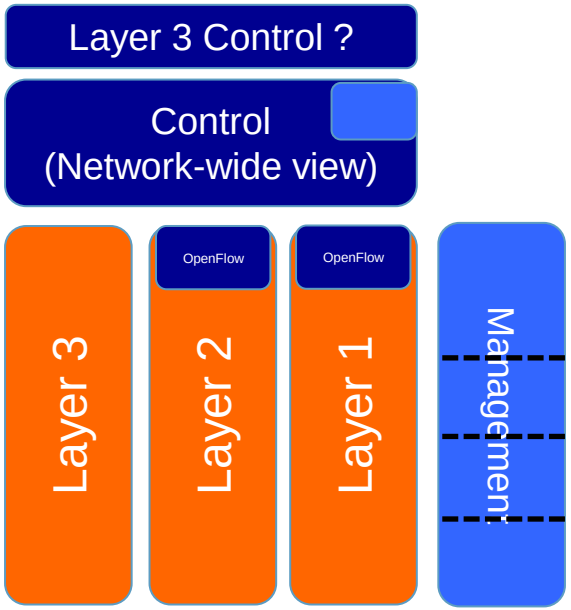

Claim #1: Programmability will lead to greater predictability

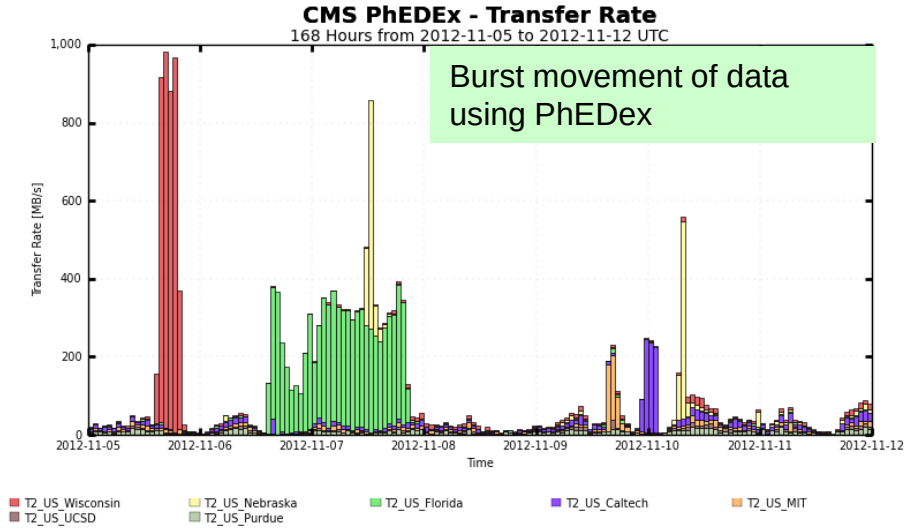

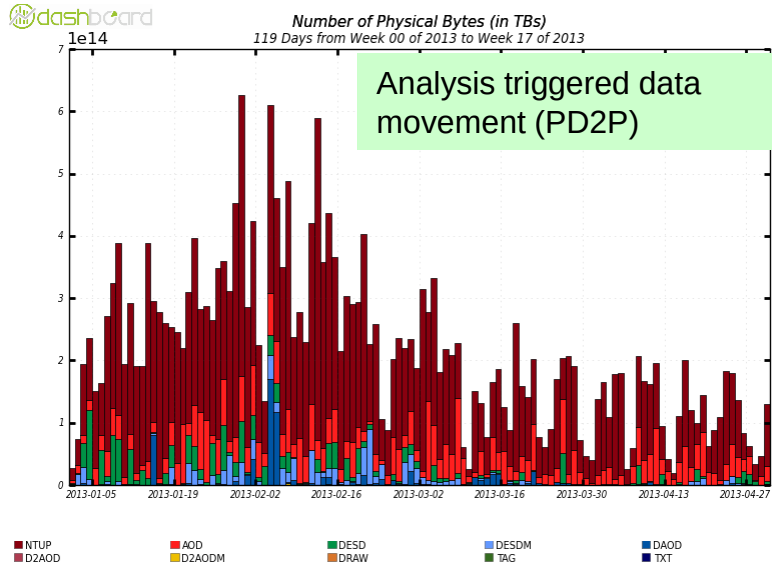

- HEP increasingly needs to deal with high performance, any-any bursts of data
- SDN enables
  - multi-layer control – packet and optical layer
  - Control over individual flows – ex. Route science flows around packet bottlenecks
- Many NRENs have access to fiber, optical and packet platforms.

ESnet leading development of SDN control for packet-optical – demonstration 10/16

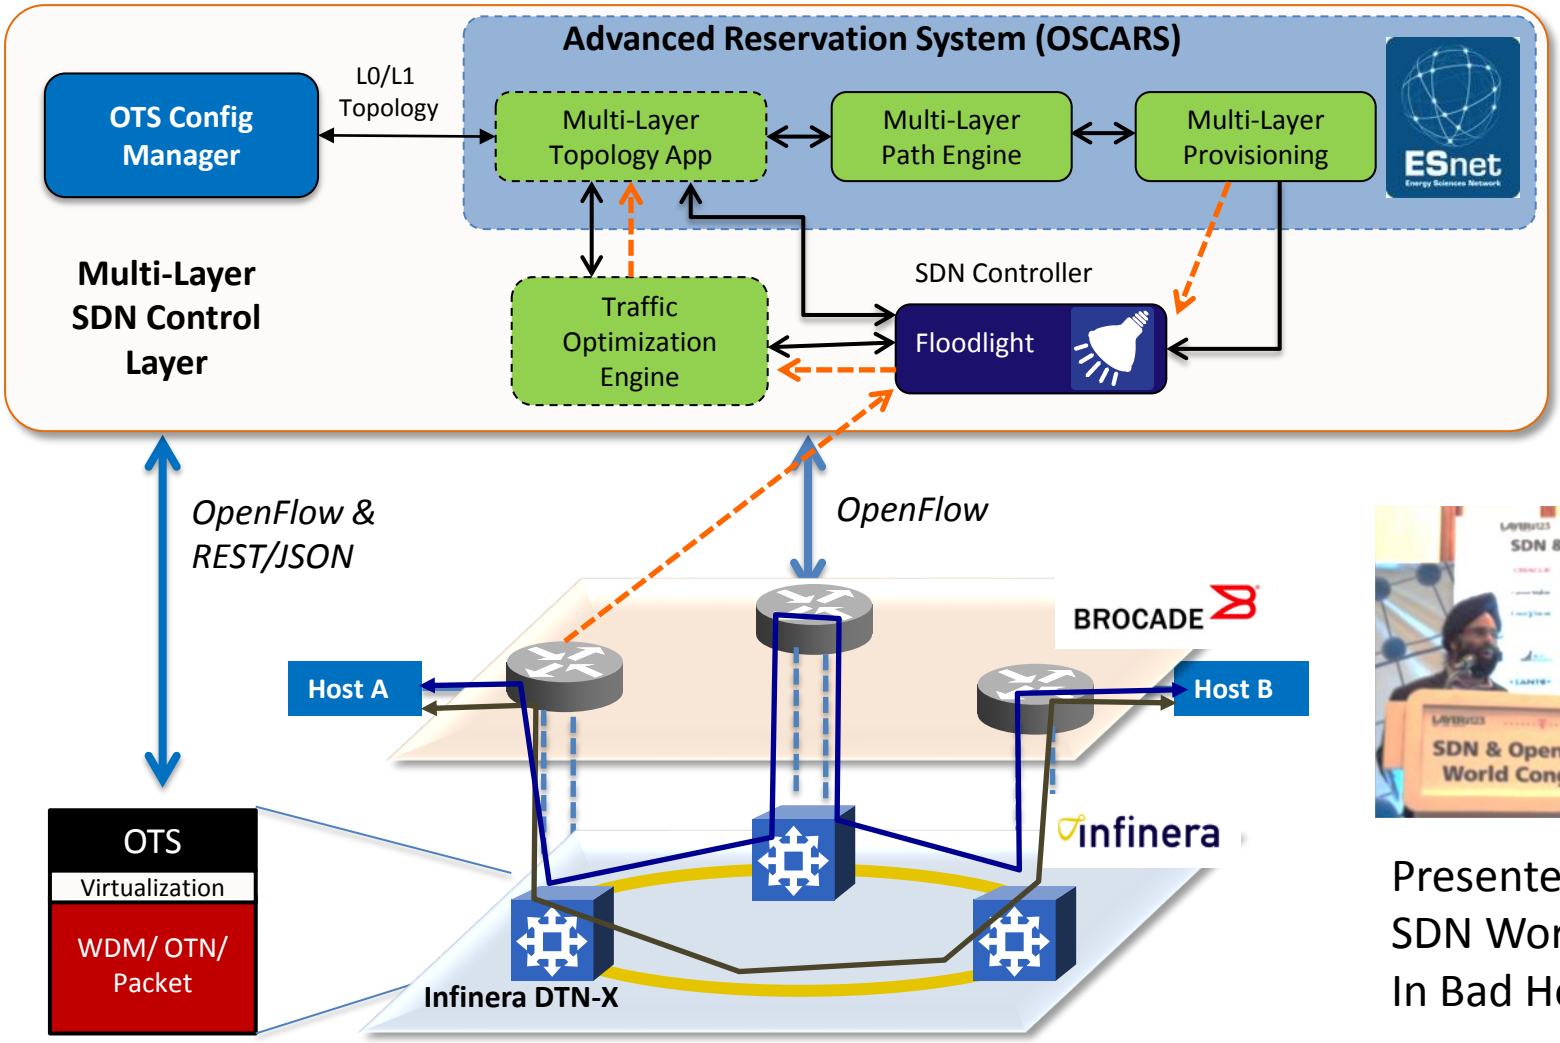

Presented at
SDN World Congress
In Bad Homburg

SDN is about network abstraction

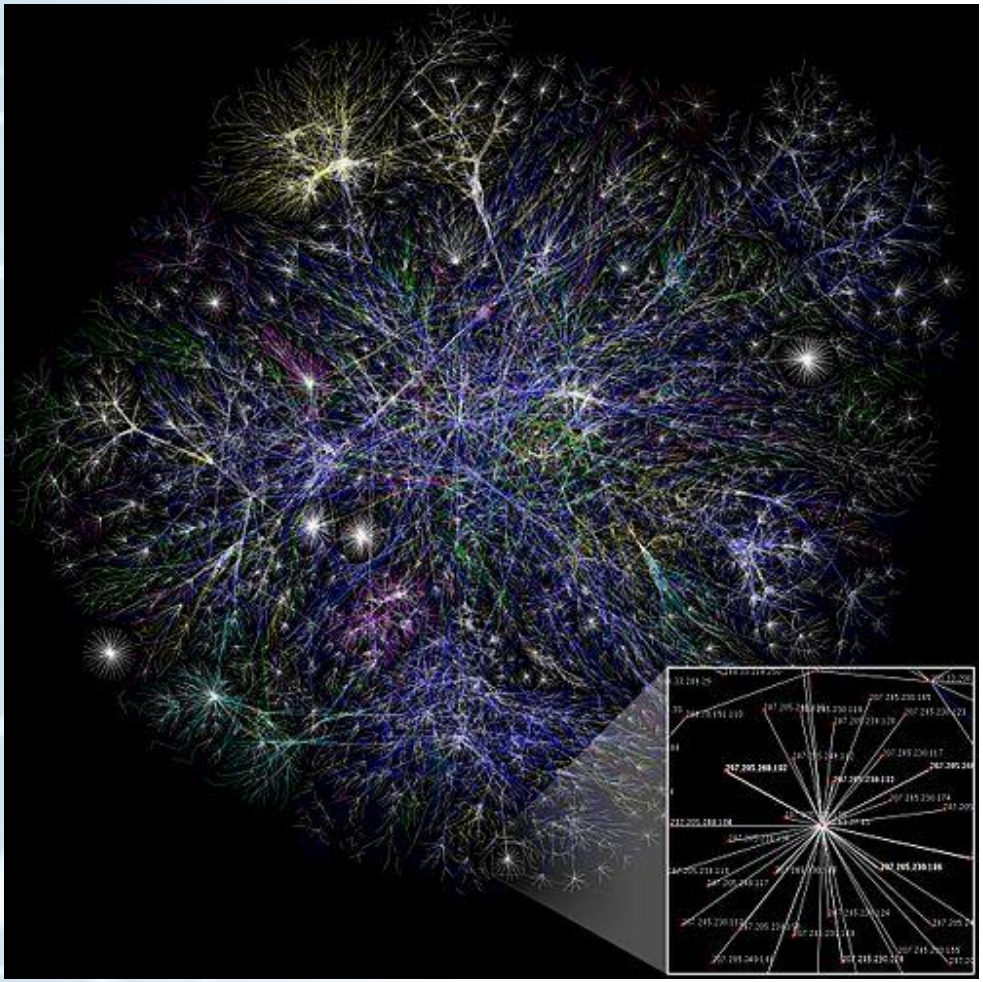

Image by Matt Britt (used with permission under Creative Commons Attribution 2.5)

Complexity of the Internet infrastructure 'black box':

It's ***strength*** and ***weakness***

Constructing simplicity through abstractions

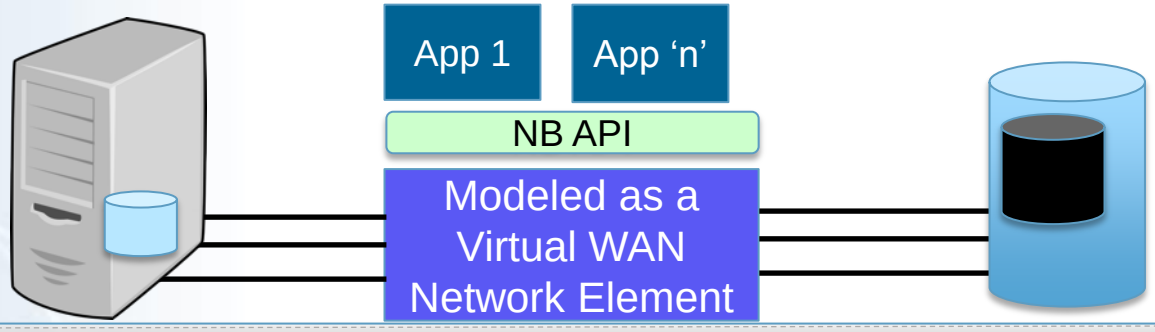

simple

complex

Network Virtualization

Network slice

Network Controller(OS)

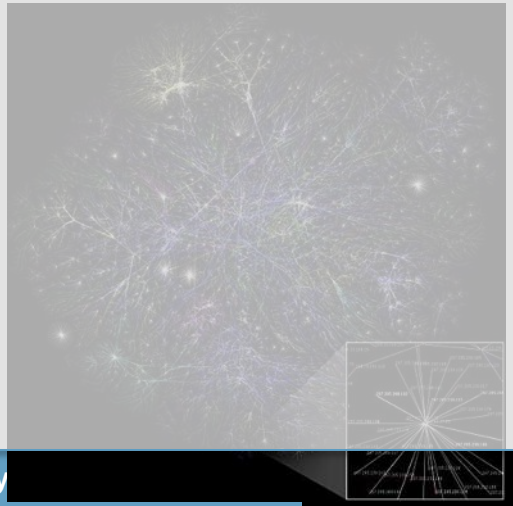

Claim #2: Virtualization will simplify how HEP applications program the network

Going from the Monarc model to full multi-point mesh would be as simple as changing the multiple point to point link abstraction with a virtual switch abstraction

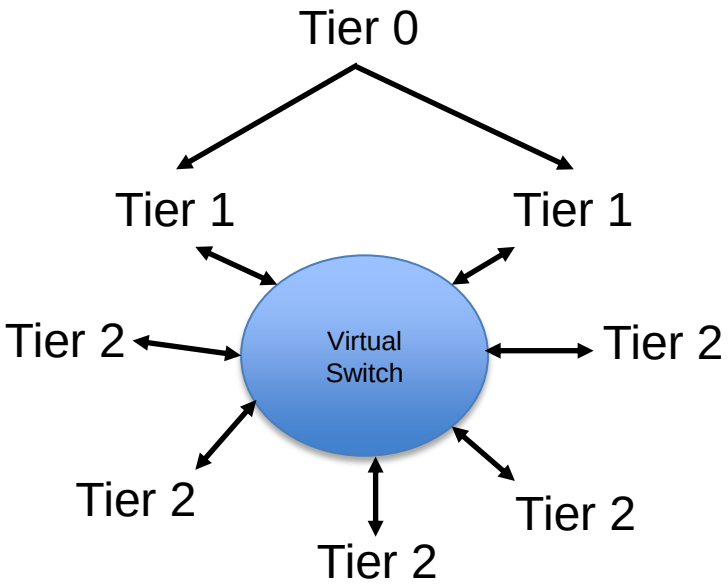

The complexity is absorbed by ‘software hypervisor’ of the underlying multi-domain network

SDN is about system optimization

Google's B4 SDN Network Utilization

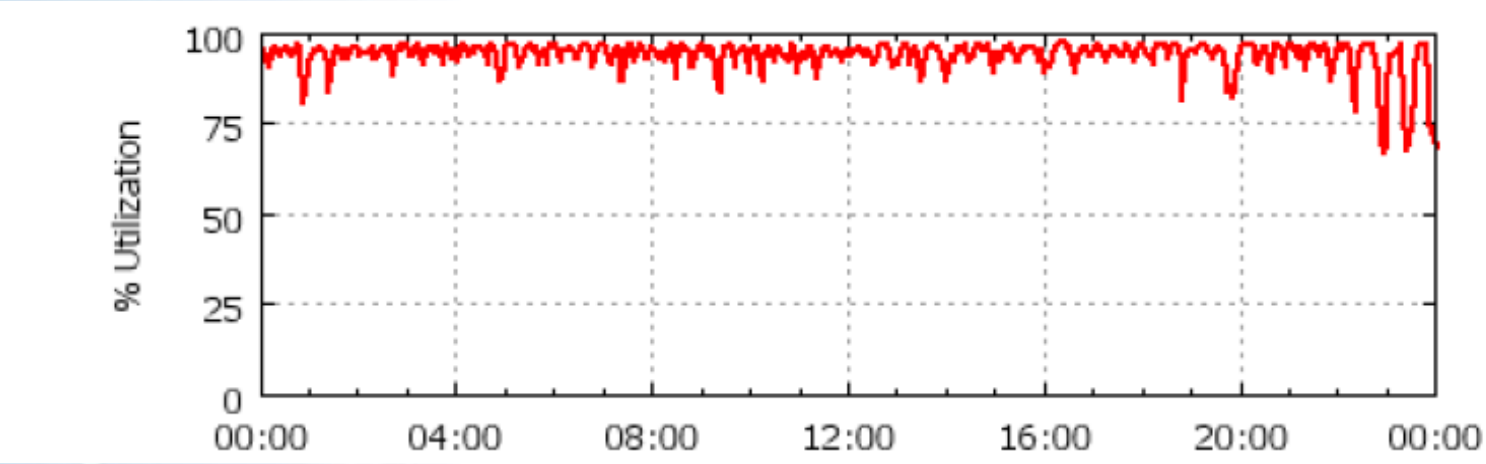

When the application and network work as a system, network resource optimization is possible

Without knowledge of flows, networking can only do coarse characterization

Fine discrimination of flows possible with SDN, meet application needs

Claim #3: SDN enables an opportunistic way to leverage all bandwidth without extra investment

exploits the fact *'In general it's much cheaper to transport data than to store it'*, for example, enables diskless Tier3

We are at the starting line!

SDN concepts have made tremendous leaps in academia and industry

There are three major leaps that still need to be made:

- Operationalization of the SDN model
- Manageability (for failures, failure recovery) and Scalability
- Security

Next Challenge is to tackle the gaps with open-source toolsets

- So a small network (not size of Google) can enjoy the potential benefits

Network organizations need to tackle similar software challenges

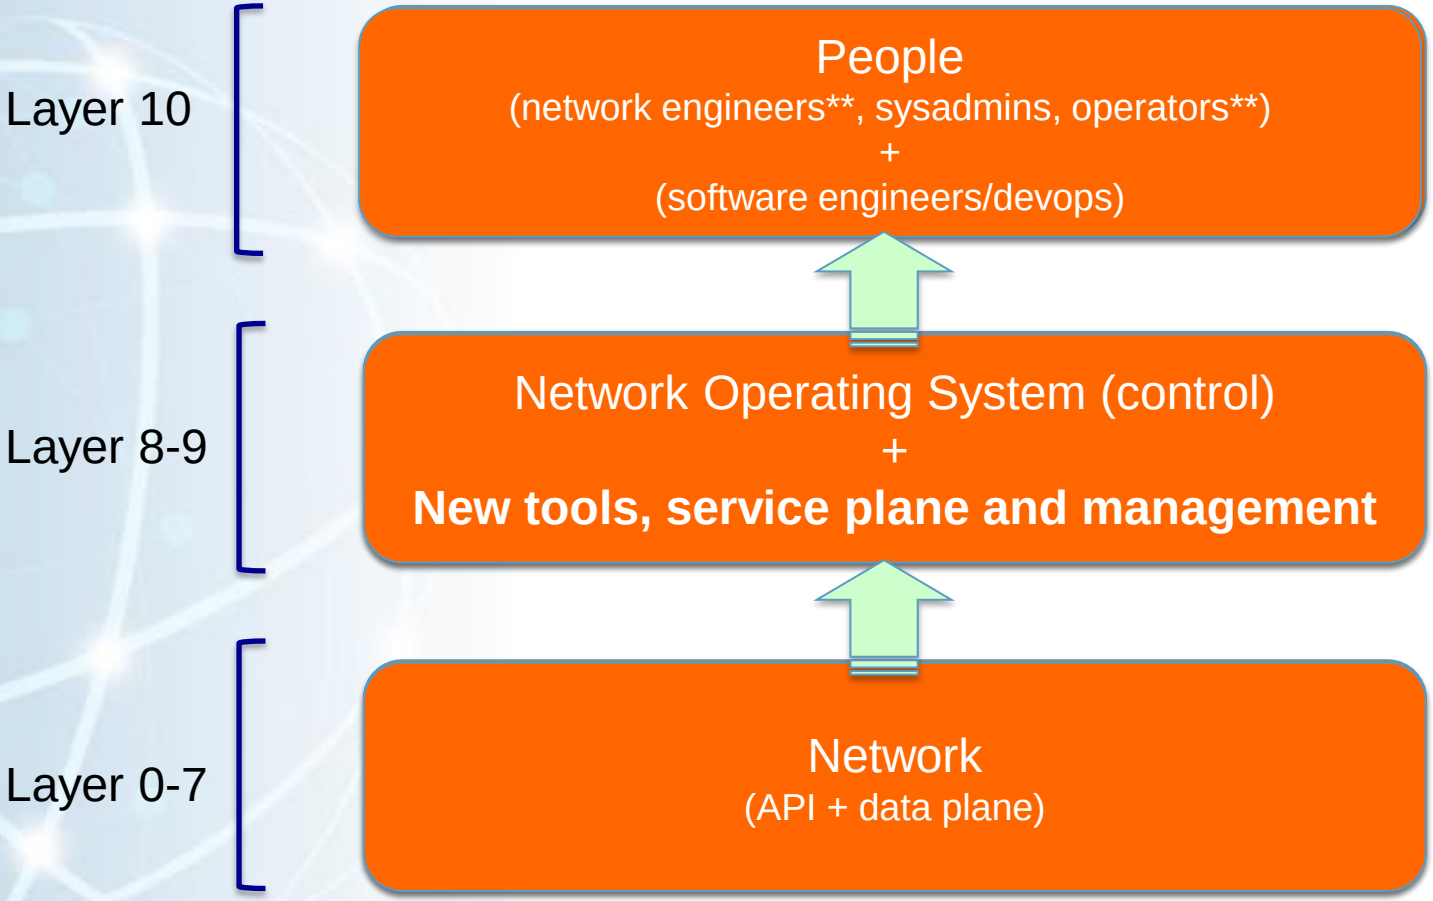

*"Living with inefficient software is much more **expensive**"*

- Torre Wenaus Keynote

** need to develop new skills

A fun peek into the future…just imagine

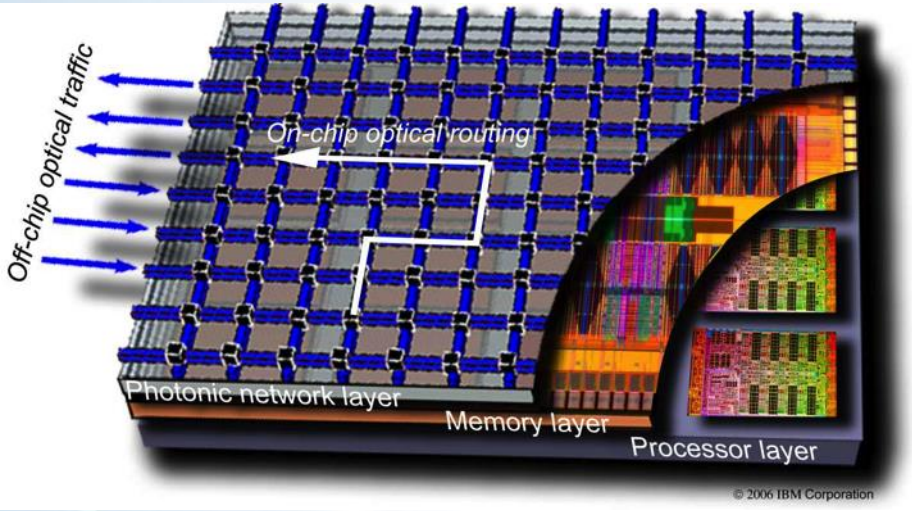

With silicon photonics integration, each chip will have a network interface

That implies each chip could be network addressable

If so, we could design servers without needing NIC cards – no difference between communication within the motherboard or outside.

With HEP applications like FAX, file systems or memory can be mounted remotely to my chip while ‘streaming data for analysis.’

With SDN, can effectively route IP and non-IP protocols (like ROCE)

SDN could revolutionize how computing is done, are we ready for that?

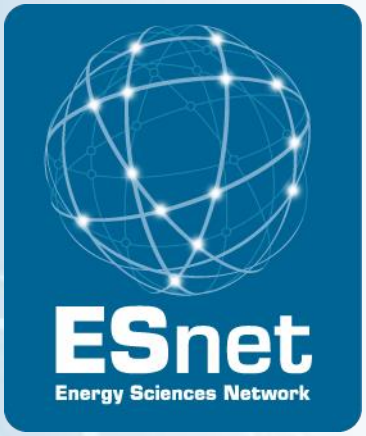

Thank you!

The wall must go!

ESnet

Researchers & Applications

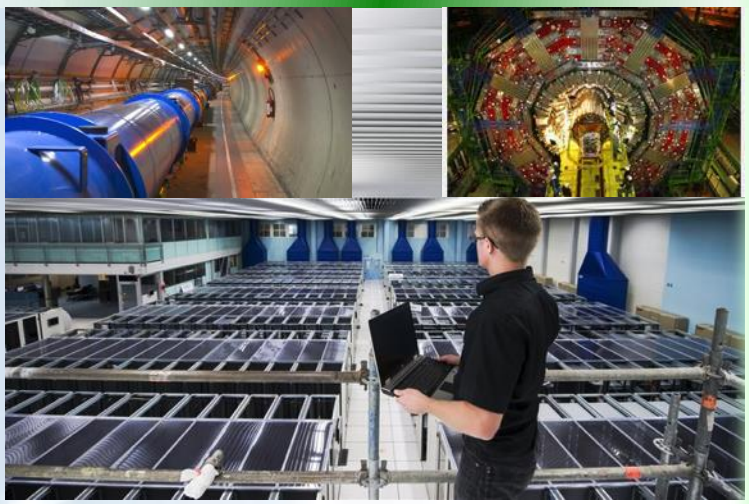

Scientific Productivity

The Network

Connections