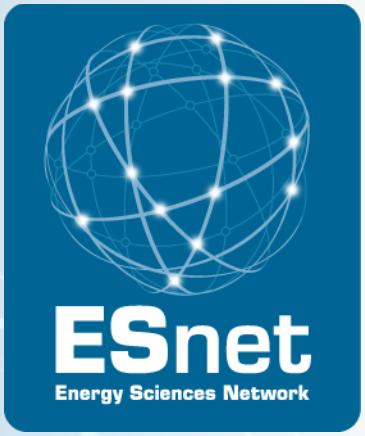

SDN for Science Networks

Inder Monga

Eric Pouyoul, Chin Guok and Eli Dart

Energy Sciences Network,

Scientific Networking Division

Disclaimer

Two Prime Requirements

1. Data Mobility

Long latencies (RTT)
Multi-domain
Multi-vendor
Multi-technology

2. Global Collaboration
   - Higgs Boson

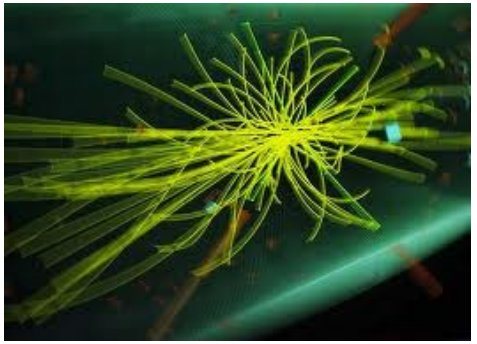

Data Mobility:
Problem is simple to articulate

Ability to move large quantities of data from one location to another over the network

a) Tuned and architected for the **best performance**

b) Most efficient utilization of **end-to-end** system capability &

c) Usability for end-user, site-admin and network-admins

Why is this a problem?

TCP is the underlying data transfer protocol

- A “fragile” workhorse
- Very sensitive to loss, especially on long RTT links
- Leads to non-ideal use of the deployed infrastructure
  - Apps underperform their potential

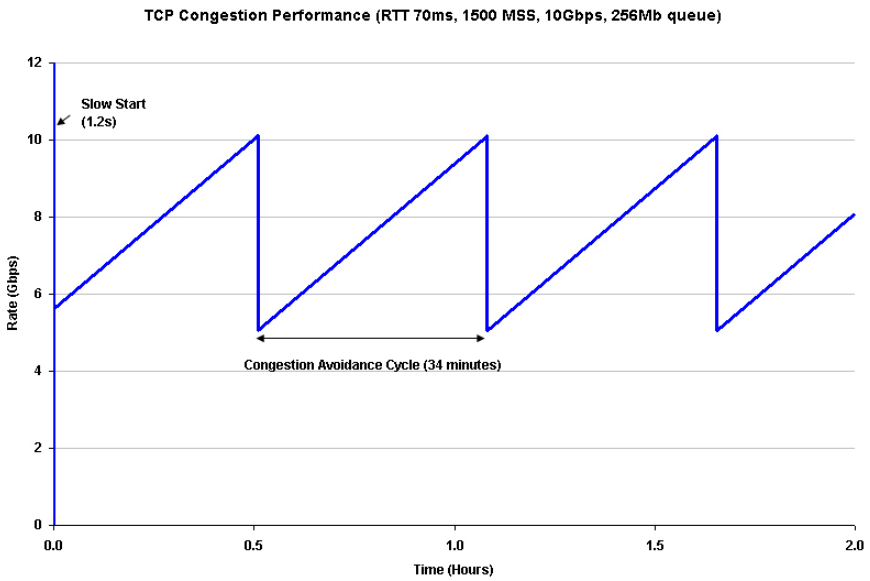

http://www.potaroo.net/papers/isoc/2005-06/fig2.jpg

Possible approaches?

1. Replace/change TCP
   - New TCP-variants or layer 2 protocols like RoCE

2. Use TCP with right environment

   Provide a loss free, high bandwidth network service over wide area
   - Enough bandwidth to avoid congestion-based loss
   - Fast lanes (virtual circuits), end-to-end
   - Big buffers to avoid burst-related loss
   - Test and measurement infrastructure to ensure a perpetually clean infrastructure

SDN/OpenFlow Investigation:

What strengths can be applied to the problems described?

Journey with OpenFlow/SDN

Joint Techs Summer 2011, Fairbanks, Alaska

Value #1

Flexible, Programmable separation of flows

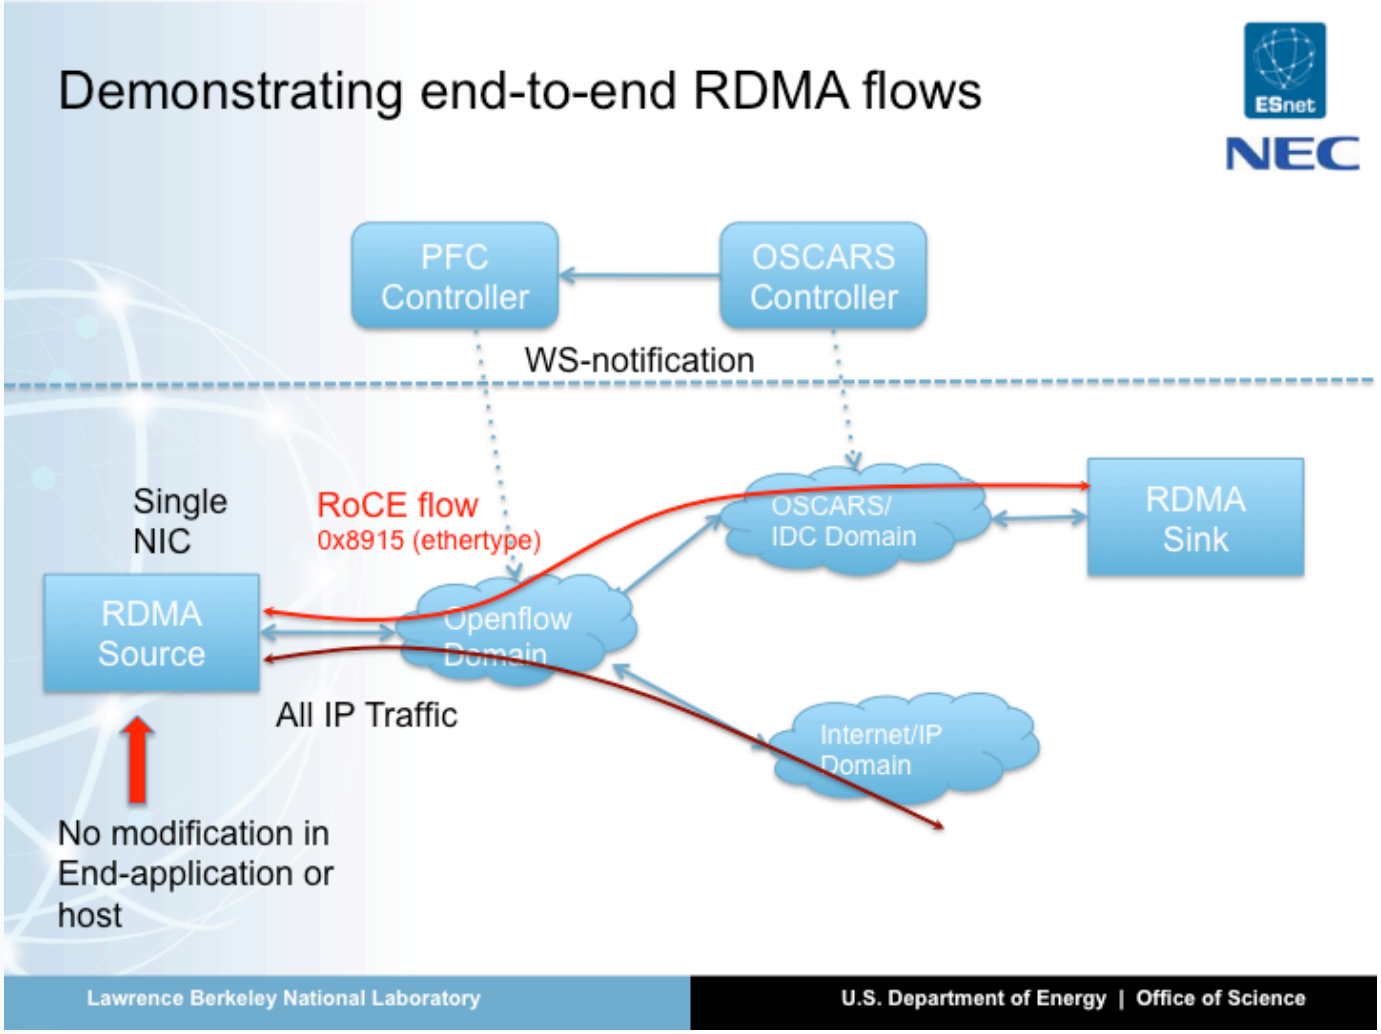

Journey with OpenFlow/SDN

Inaugural Open Network Summit, 2011, Stanford and SuperComputing 2011, Seattle

<u>Value #2</u>

End-to-end conversation at each layer is important for stitching

(even at layer 8)

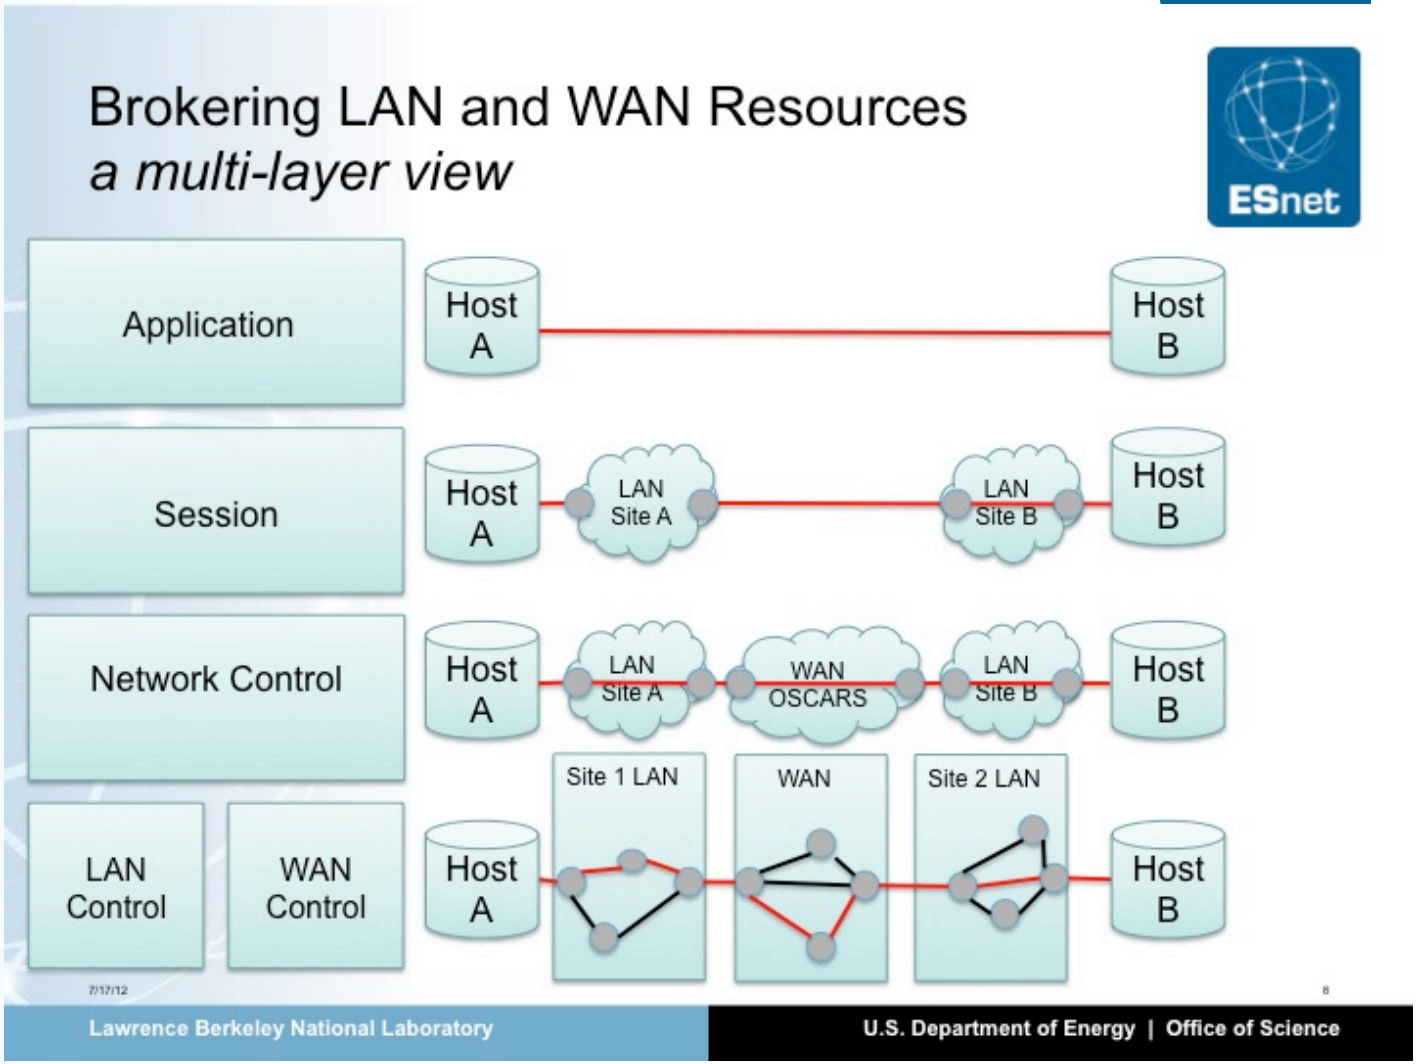

The “ScienceDMZ” Design Pattern

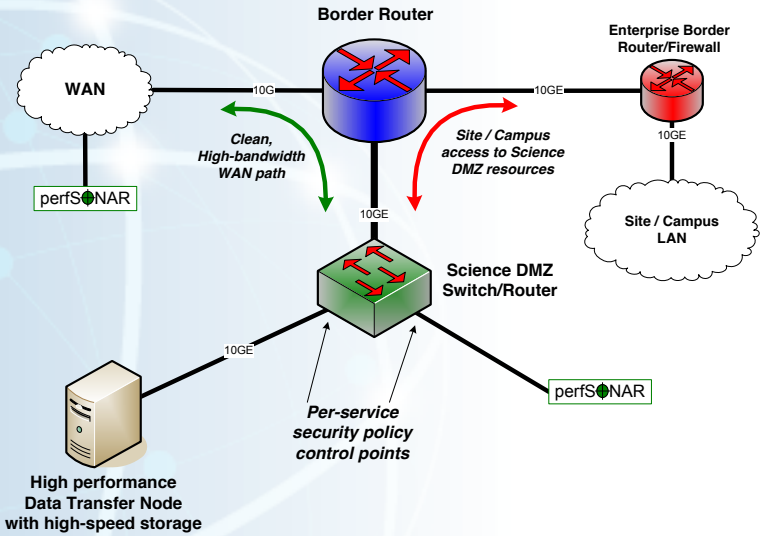

Vanilla Science DMZ

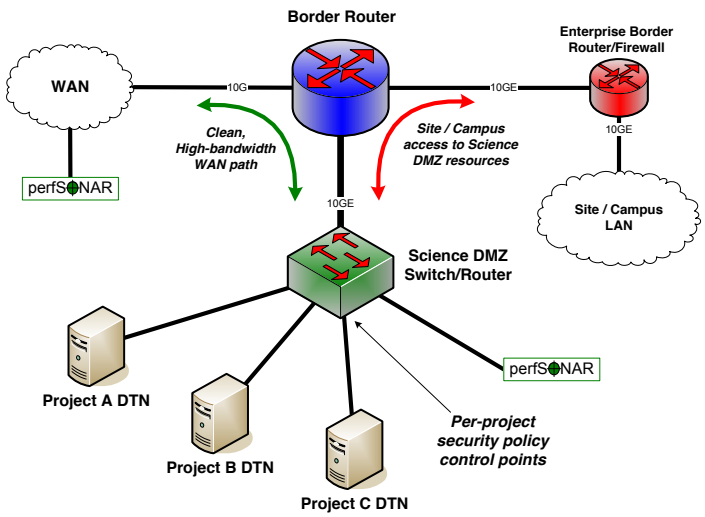

Multi-Science DMZ

Desired capabilities

- Automated end-to-end connections (less dependent on real-time human decisions)
- Best utilization of shared WAN resources
- Dynamic – flexible based on application usage patterns

Implementing the ScienceDMZ Design Pattern

Simple application of OF

OF Switch: fine grained mapping of science flows to guaranteed bandwidth circuits

- Dynamic
- Application/Policy driven
- Automated VLAN translation

OF Controller: manage WAN resources (virtual circuits, bandwidth etc.)

- Site administrative resource allocation
- Site-WAN, Site-Site policies enforced

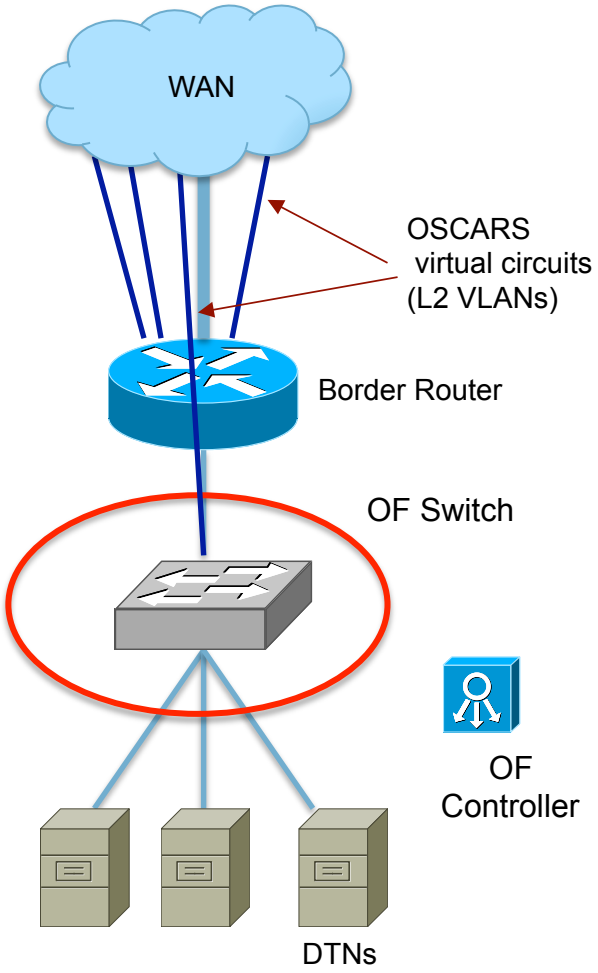

How about the PIs with their storage and clusters?

- Firewall Data Planes should not be involved in science data flows
- Firewall rule-sets for science services become self-maintaining
  - i.e. the rule-set tracks service provisioning

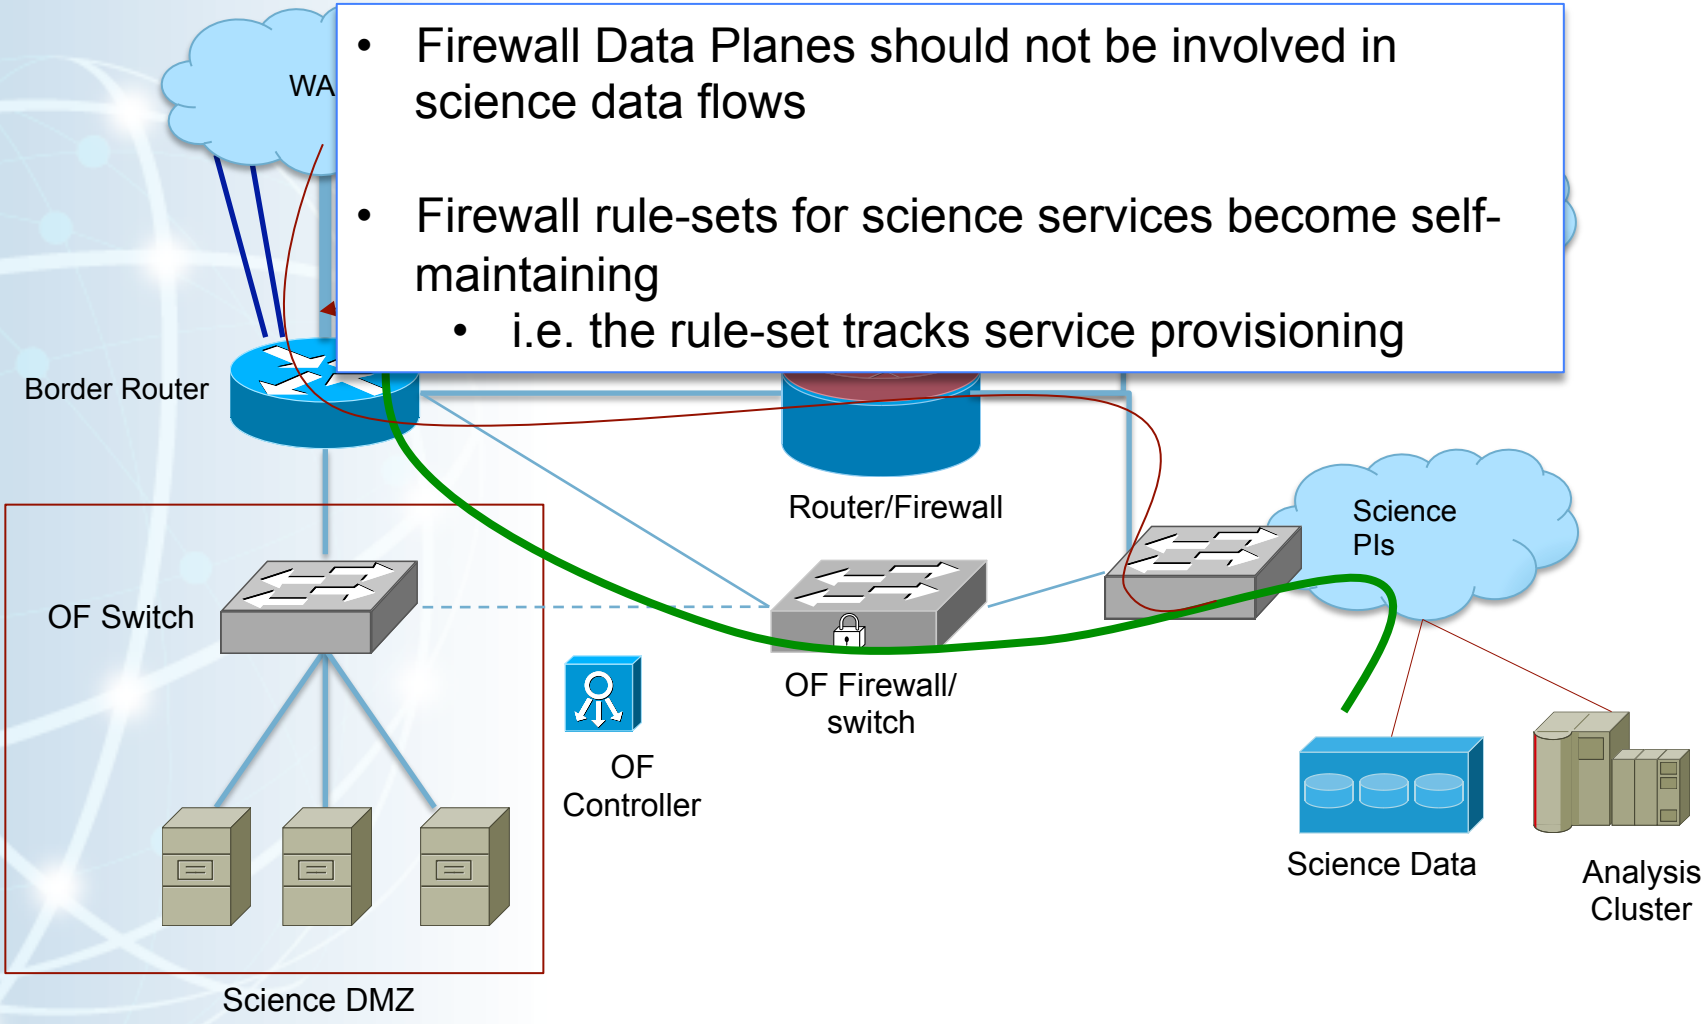

SDN and the Wide-Area Network

Software-Defined Networking has already been well adopted by the R&E wide-area networks

- OSCARS centralized advanced reservation and provisioning

What's different with this SDN/OF wave?

- Formal concept of a network OS
- Abstract Network model

What are the fundamental network abstractions?

- All discussions on standard Northbound APIs are fruitless unless we define these

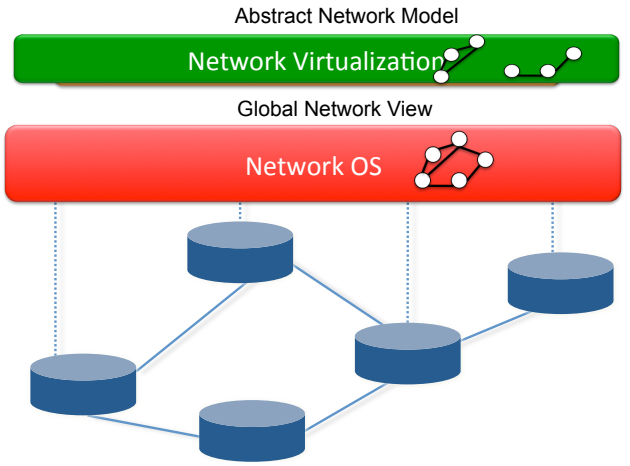

SDN with Virtualization
Guru Parulkar, from
"OpenFlow/SDN: A new approach to networking"

Going back to Science Networking requirements

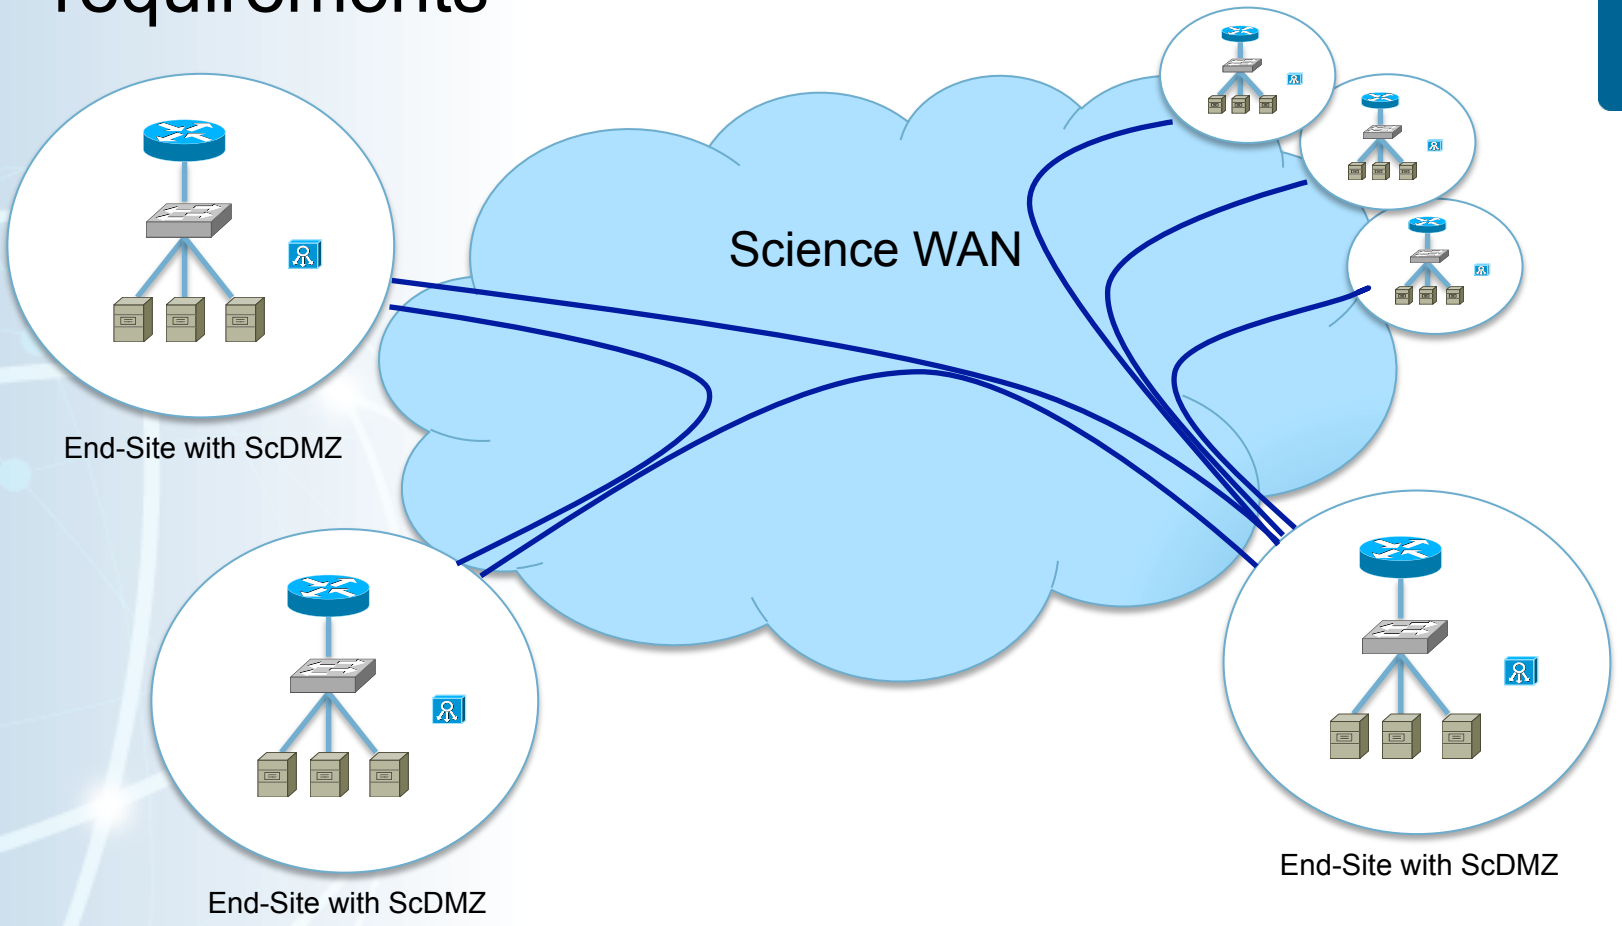

Dynamic Point to point circuits scale reasonable well,
but don't meet all global collaboration requirements (requirement #2)

A wide-area abstraction = Logical programmable OF switch

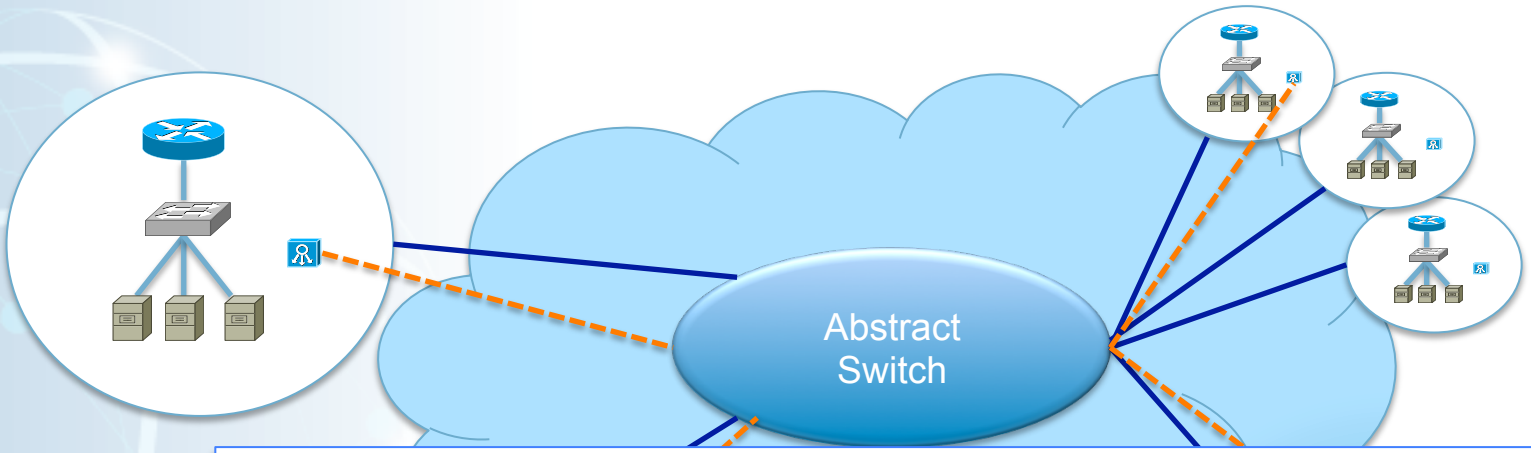

- Multi-point to multi-point connectivity
  - While leveraging the multi-domain, advanced reservation capabilities of R&E networks
- OpenFlow interface for flow programmability by the ScDMZ OF controllers
- Can be sliced further into virtual-switches or topologies

Recursive Abstraction

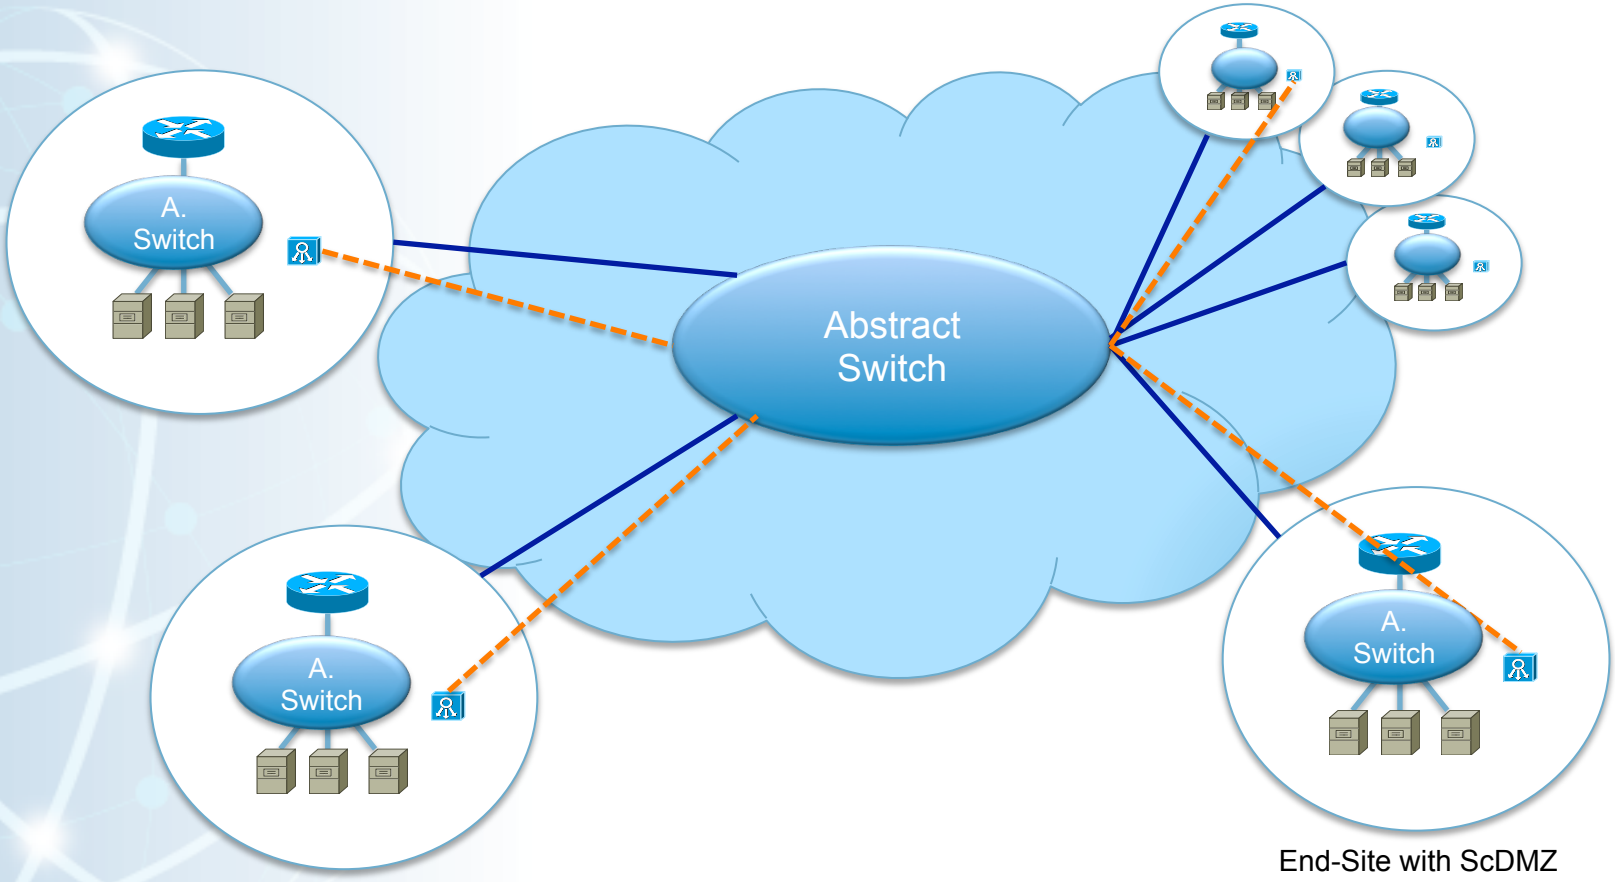

Practical considerations of a programmable switch abstraction

- Do not need to have all OF devices in the WAN
- Do not need to have OF support in the Site
  - Just a controller
- No new protocols or API
- Capable of supporting both L2/L3 switching
- Supports all models of end-to-end conversations aka brokers
  - ECSEL, GENI…

Summary

Mapping Science requirements, architectures, design patterns to OpenFlow/SDN paradigm

- Still in exploration phase
- Looking at models that does not force a complete overhaul of the network

Simple abstractions are needed, that scale from campus to WAN

Layer 8 issues are important and will be hotly debated

- Automation and policies go hand in hand