



ESnet

ENERGY SCIENCES NETWORK

ESnet's (100G) SDN Testbed

Inder Monga and ESnet SDN team

International SDN Testbed,
March 2015



U.S. DEPARTMENT OF
ENERGY
Office of Science

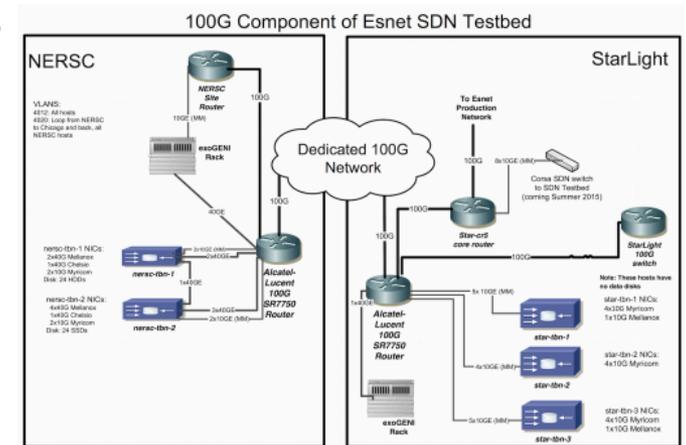


Outline

- Testbeds in ESnet
- Motivation: Building a Scalable SDN WAN testbed
- Hardware and Deployment Status
- First Experiments

Testbeds in ESnet

- Operating 100G Testbed and multi-layer testbed since 2011
- Testbed Model = 2-3 page application to use, calendar based access, yearly renewal
- Lots of published research, work on TCP, flow classification, 100G IDS etc.
- 100G now common service to many campuses, and in many networks
- Migration to 100G SDN Testbed!



Building a Scalable SDN WAN Testbed

1. Prototype for a potential ESnet6 design
2. Multi-domain federation for the science complex
3. International SDN collaboration

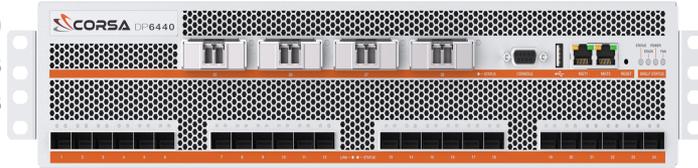
Many concepts work well in lab, but require significant modification to run @ scale

Building a Scalable SDN WAN Testbed

- Assumption: Greenfield SDN testbeds need to interoperate with other non-SDN networks or may need to federate over existing non-SDN networks
 - Isolation is **NOT** a feature we want on this testbed
- Requirements
 - L2 / L3 transport with QoS, MPLS
 - Support for federation
 - Internet-scale routing and forwarding tables
 - OpenFlow 1.3 support
 - Diagnostic and troubleshooting capabilities

Testbed Hardware: Corsa Technology

DP6440
4 x 100G CFP2 Ports
24x 10G SFP+ Ports

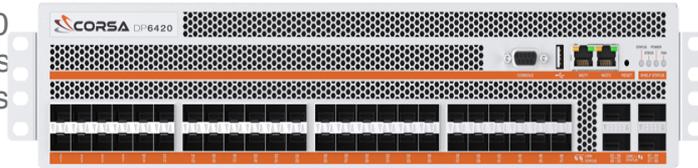


DP6430
2 x 100G CFP2 Ports
24x 10G SFP+ Ports



100G

DP6420
48 x 10G SFP+ Ports
4 x 40G QSFP+ Ports



10G

DP6410
24 x 10G SFP+ Ports



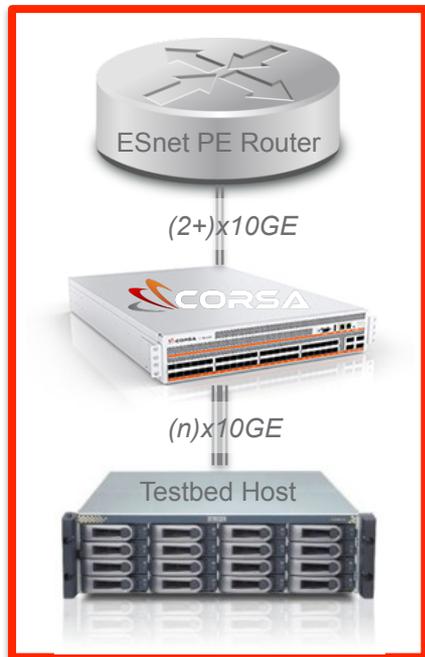
Open vSwitch + 
OpenFlow 1.3 / 1.4 / 1.5+

- Fully programmable OpenFlow data plane with multiple pipelines
- Embedded OVS based OpenFlow 1.3+ / ovsdb SDN controller interface
- Ultra fast flow setup and control plane interfaces
- Full line rate, OpenFlow ports with multiple flow tables on all ports.

Notable Features of Corsa Switches

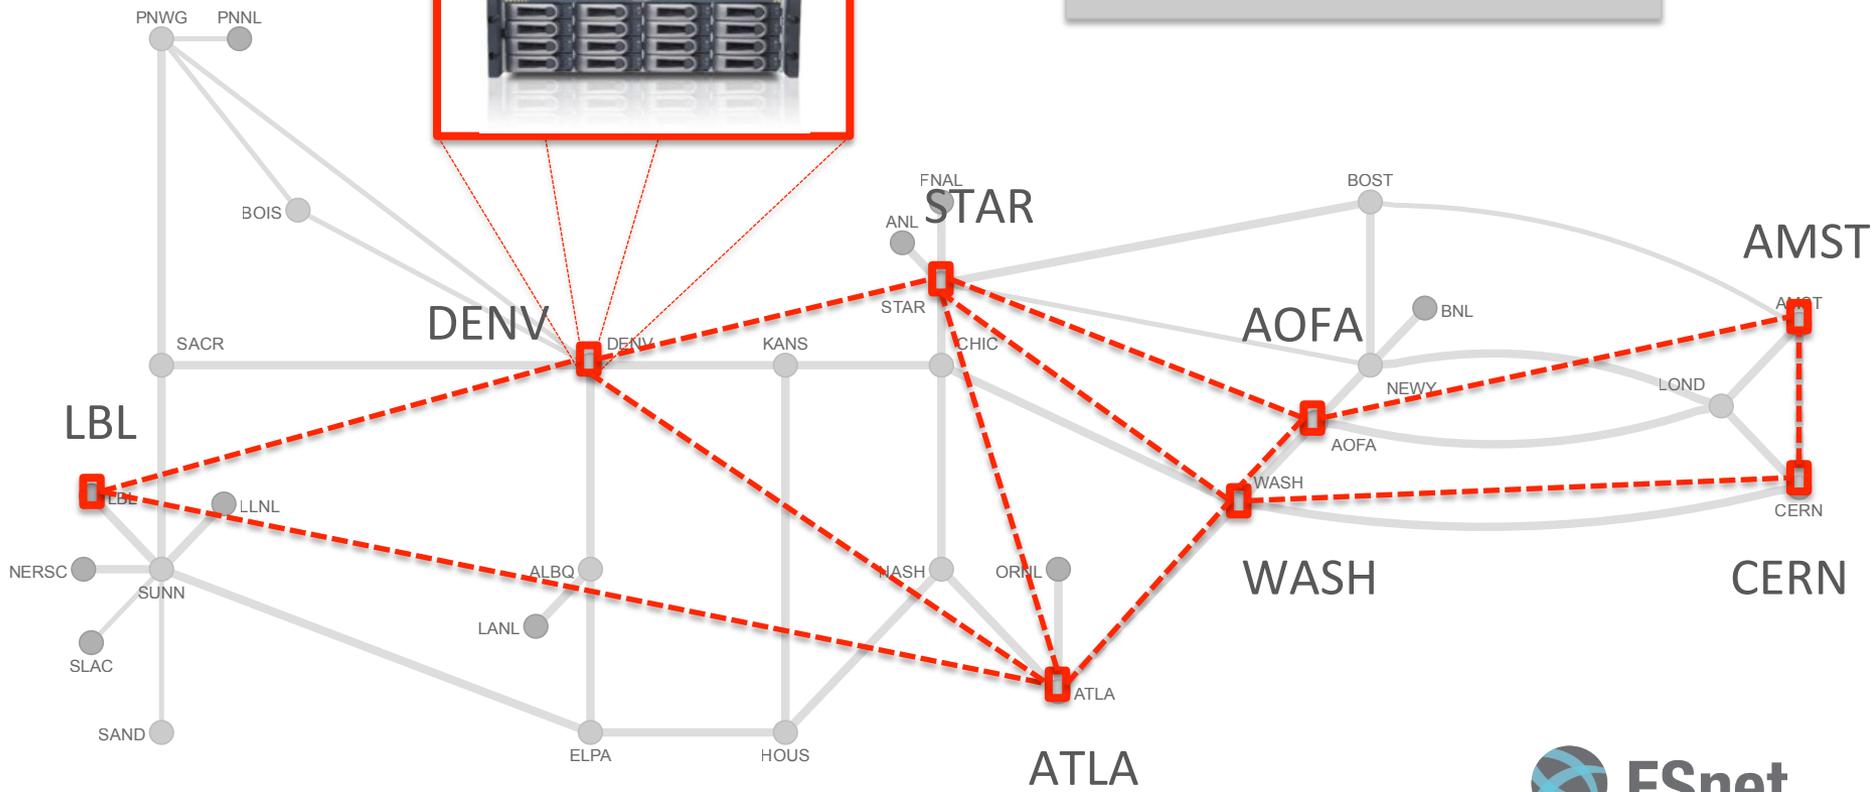
1. Multiple flow tables
2. Millions of flow entries
3. Large scale packet buffers (20GB)
4. 10- and 100-Gbps ports at line rate
5. Extremely fast flow modifications: 10K—1M entries per second (as opposed to 100 flow mods / sec)
6. OpenFlow 1.3+ protocol support

ESnet SDN Testbed

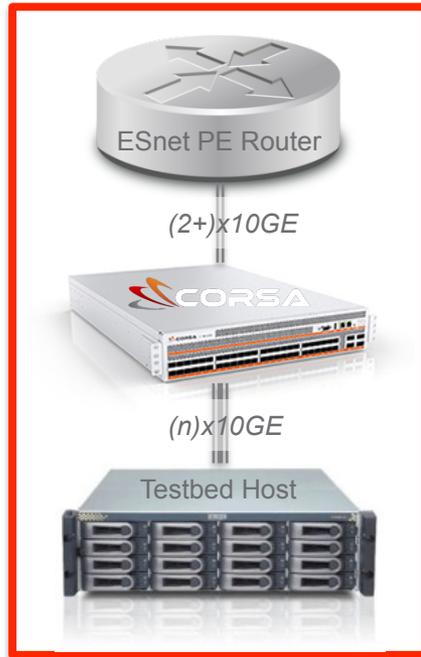


 Planned SDN Testbed node locations

 Planned SDN Testbed connectivity overlay (using OSCARS circuits)



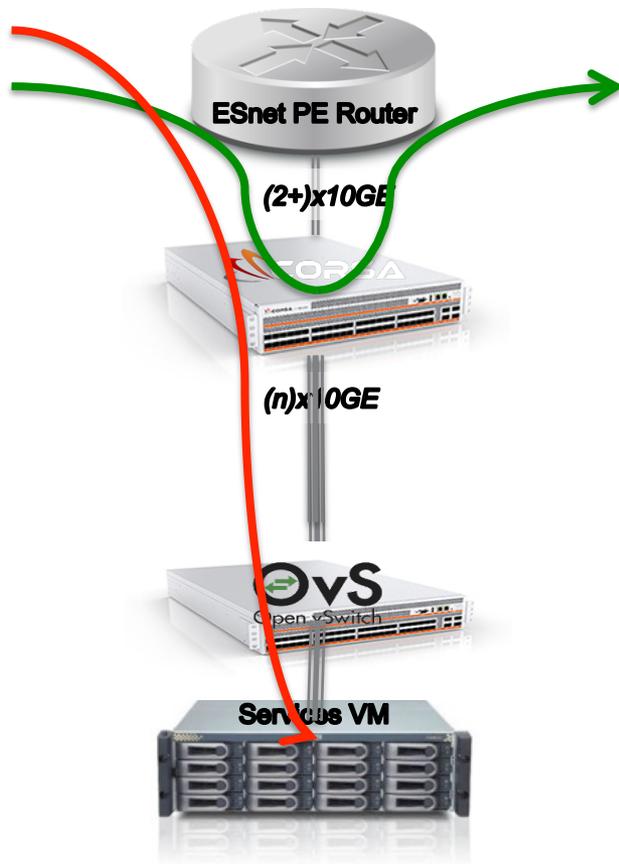
ESnet SDN Testbed Node Physical View



ESnet SDN Testbed Node Logical View



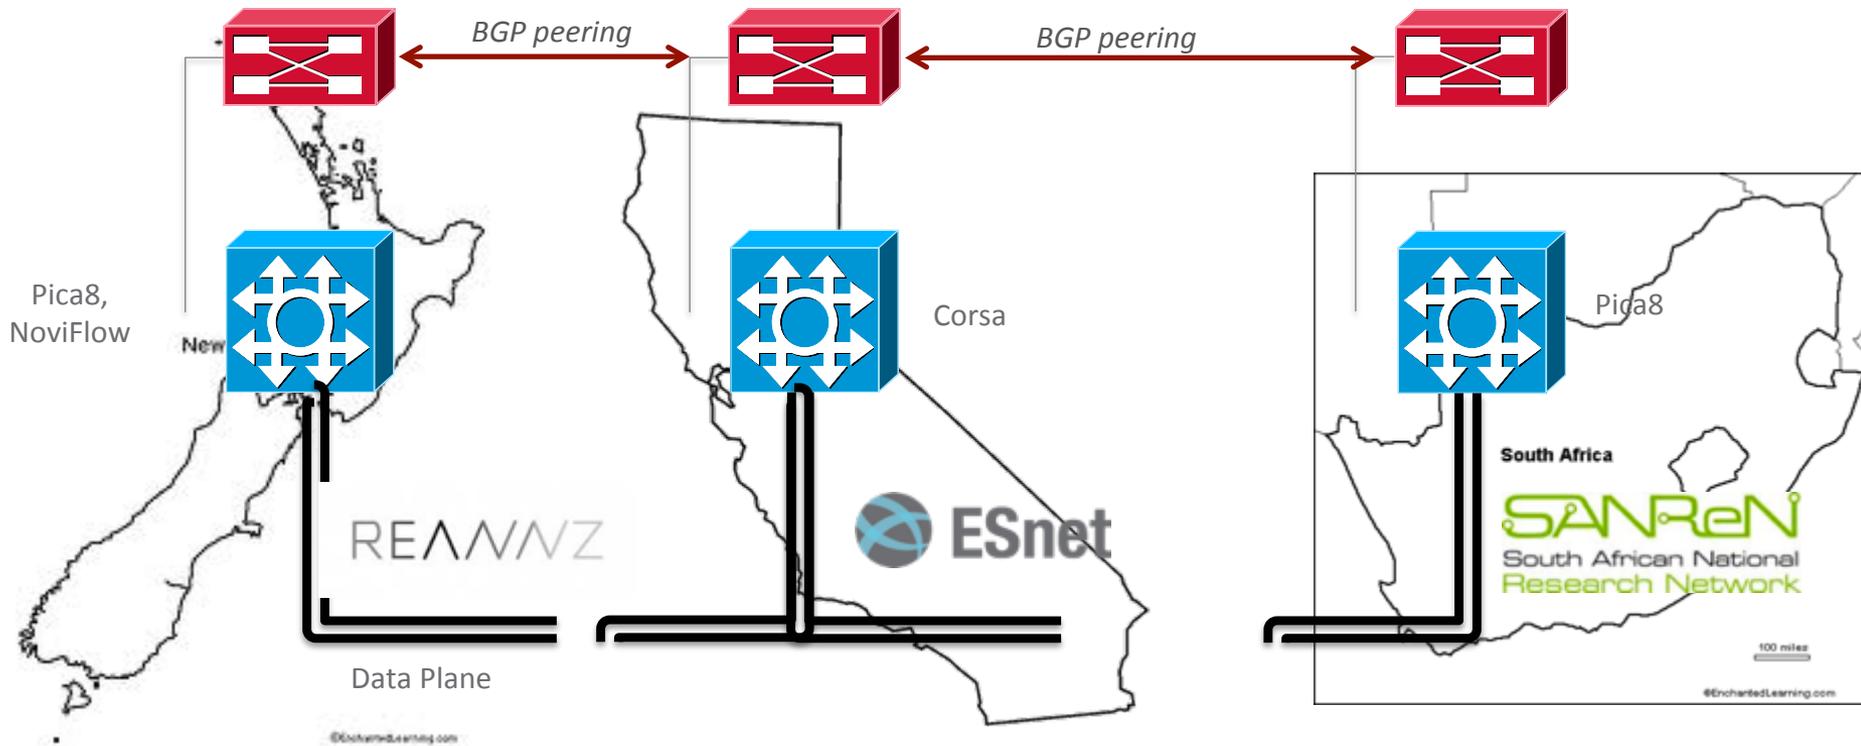
ESnet SDN Testbed Node Logical View



- Software switch and Services VM paired with every hardware switch.
- Most flows only pass through hardware switch
- Flows requiring special handling go through software switch, possibly to Services VM.
- Hardware switch provides performance, stability
- Software switch provides flexibility, without compromising stability of network

ESnet SDX: Multi-continent SDN BGP peering

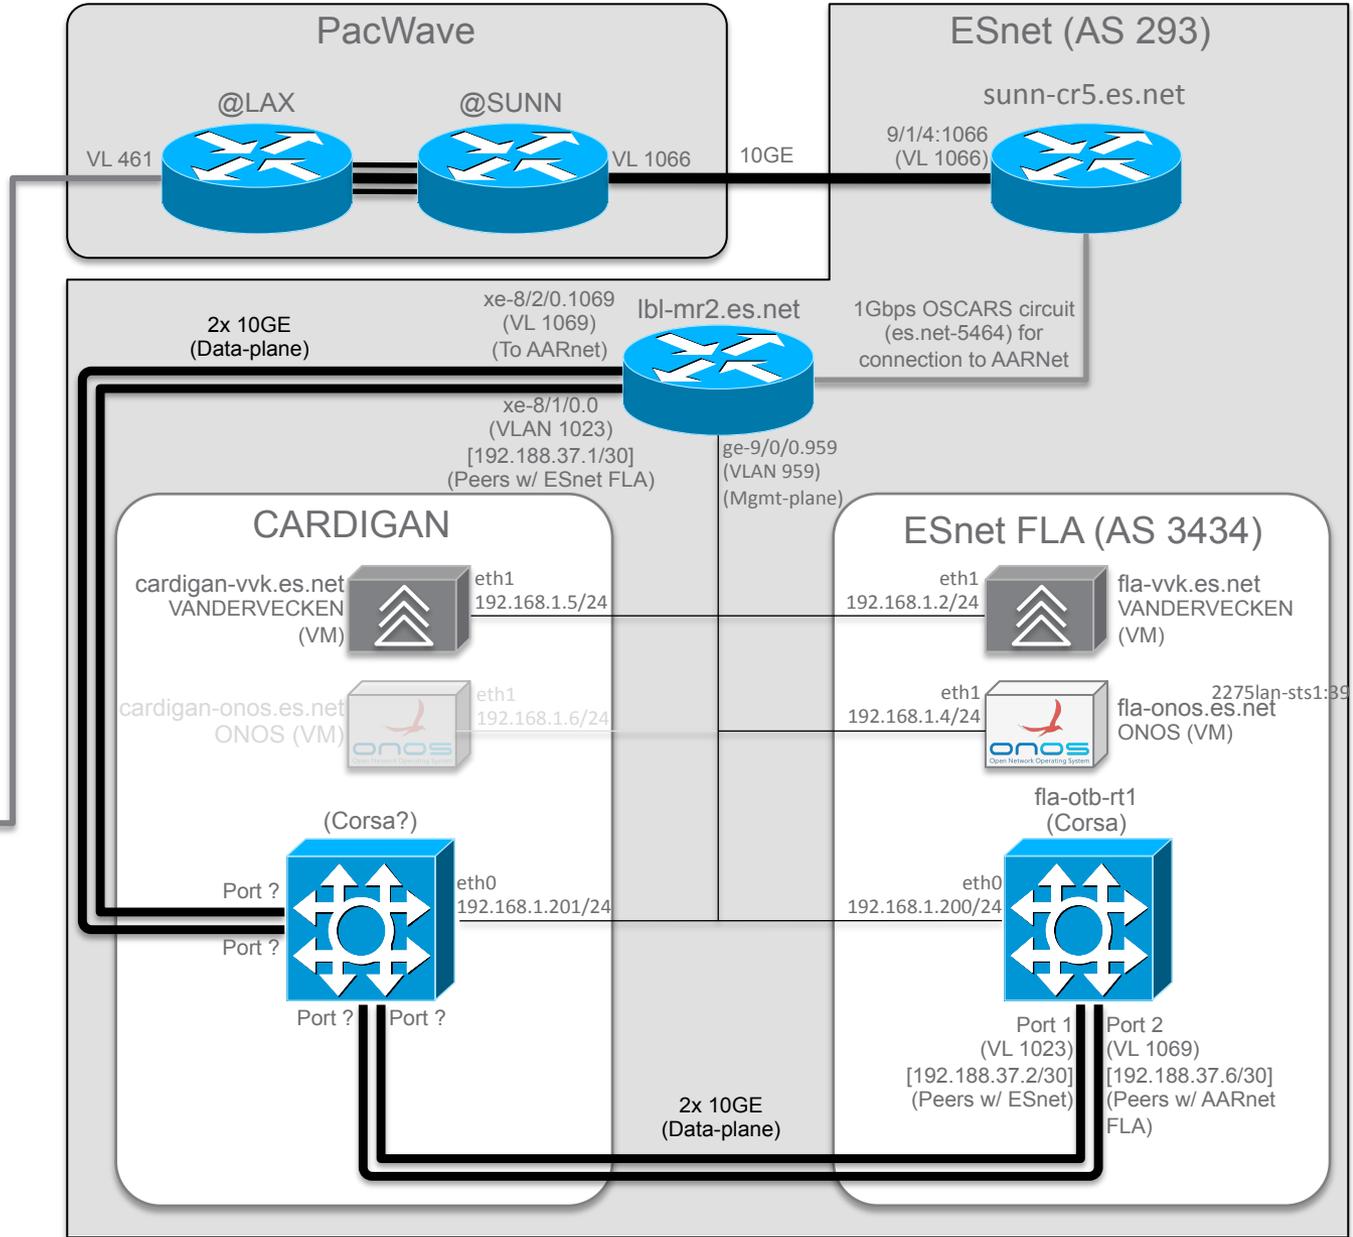
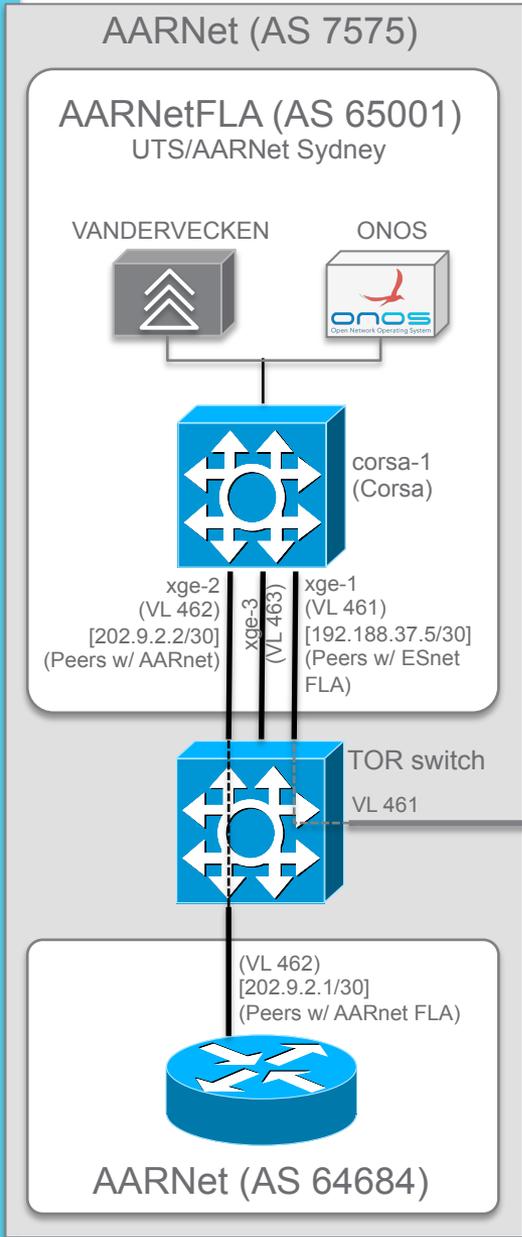
VANDERVECKEN (VM):
SDN Router



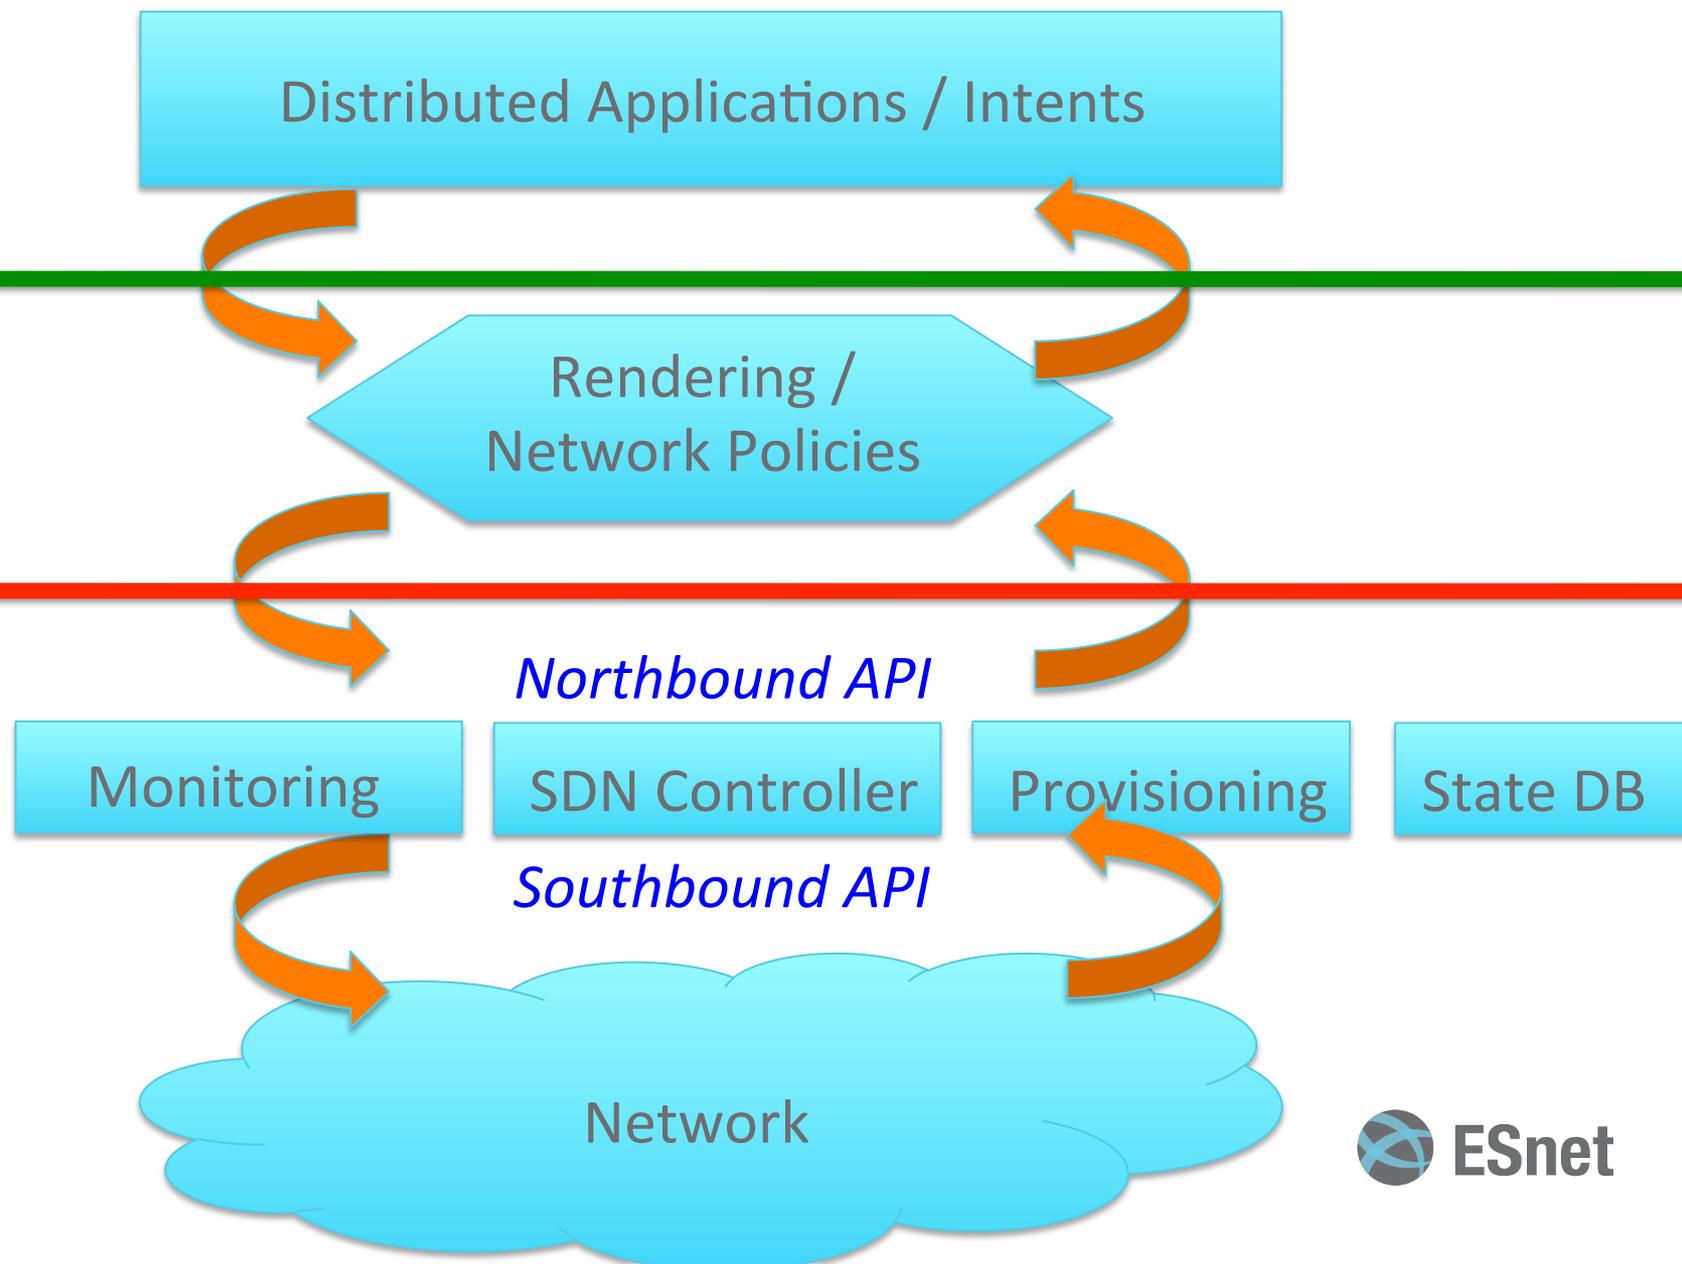
With help from -



AARNet/ESnet TreeHouse Setup [Jan, 2015]



Network Policies and Framework to support multi-tenant science applications



ENOS: ESnet Network Operating System

Goal: Provide an environment for network policies to run

- Allows applications / users to:
 - Provision endpoints
 - Specify dynamic QoS
- Allows ESnet to easily do:
 - Automation
 - Analysis / verification
- Provides a "network shell" interface to advanced network services such as:
 - Show the path options between two sites and their performance characteristics
 - Send traffic to or from some subnet to a IDS for analysis
 - Creating a multi-point VPN between sites



ESnet SDN Testbed: Current Status

- SDN Testbed
 - Testbed hardware deployment recently started
 - Plan for everything to be deployed by the end of May
- SDX
 - Initial demonstration in ONS 2013
 - Latest topology demonstrated and announced in November 2014
 - Currently peering with AARnet
- ENOS
 - Code being developed
 - Multipoint VPN being tested on Mininet

Questions?

imonga, chin, lomax, bmah, bltierney at es dot net

Thank you!

