AmLight ExP
Americas Lightpaths Express & Protect

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Supporting experiments with SDN on AmLight and AtlanticWave-SDX

SDN Optimized Advanced Network Services for Experiments
May 17, 2016

Julio Ibarra, PI
Heidi Morgan, Co-PI
Donald Cox, Co-PI
Jeronimo Bezerra, Chief Network Engineer

U.S.-Latin America 100G Link

PRESS RELEASE
FOR IMMEDIATE RELEASE
Contact: Liz Boten, eboten@internet2.edu

Americas Lightpaths Express and Protect Activates First US – Latin America 100G Networking Link Enhancing Infrastructure for Research and Education

Miami, Florida, May 11, 2016 – Florida International University's Center for Internet Augmented Research and Assessment (CIARA) is pleased to announce the first 100G research and education network link between the U.S. and Latin America, a major component of the five-year AmLight ExP (IRNC BACKBONE: Americas Lightpaths Express and Protect) NSF Award#ACI-1451018.

On April 18, 2016 the AmLight Consortium activated the first 100G link of the AmLight-ExP project. It has 106ms delay and it goes via the Atlantic between Miami, FL and Sao Paulo, Brazil. The 100G link is under evaluation for the next 30 days. "To date, we have not seen any packet loss or errors and, to evaluate it, we are using an IXIA 100G packet generator," said Jeronimo Bezerra AmLight Chief Network Engineer.

The AmLight Consortium is a group of not-for-profit universities, state, national and regional research and education networks including the AmLight ExP project at Florida International University, RNP, ANSP, RedClara, REUNA, FLR, AURA, Latin American Nautilus, and Internet2.

https://www.internet2.edu/news/detail/10882/

AmLight-ExP Today

- NSF International Research Network Connections, award# ACI-1451018, U.S.-Latin American connectivity
- 100G Miami-São Paulo, Atlantic
- 100G Miami-São Paulo, Pacific, by early June
- 6x10G links, landings in São Paulo, Fortaleza, Santiago
- 160G of aggregate bandwidth capacity

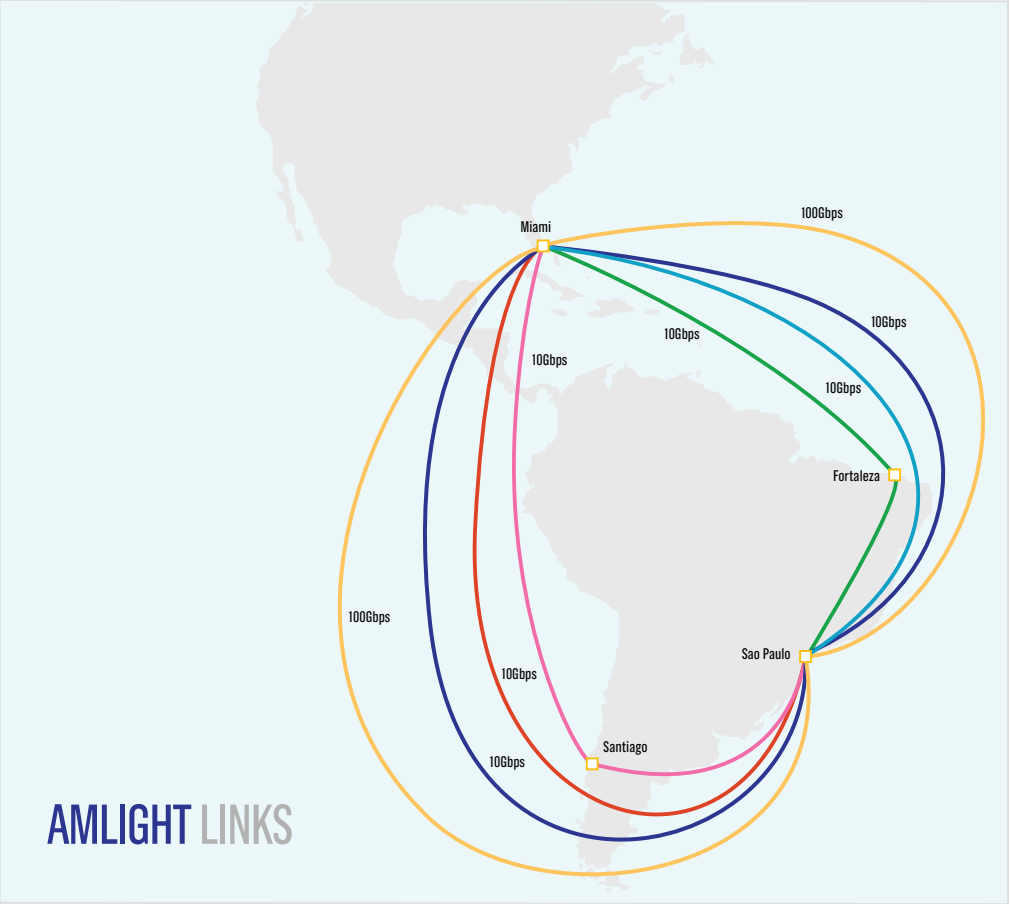

AmLight-ExP Future

- AmLight ExP 2017:
  - 600G of aggregate bandwidth to be added:
    - Santiago-São Paulo, and São Paulo-Miami
  - 100G for Panama in 2018
- Fortaleza as a south Atlantic hub
  - EulaLink submarine cable from Fortaleza to Portugal
  - SACS submarine cable to Angola (Q3 2018)
  - CBCS submarine cable to Cameroon (Q4 2017)

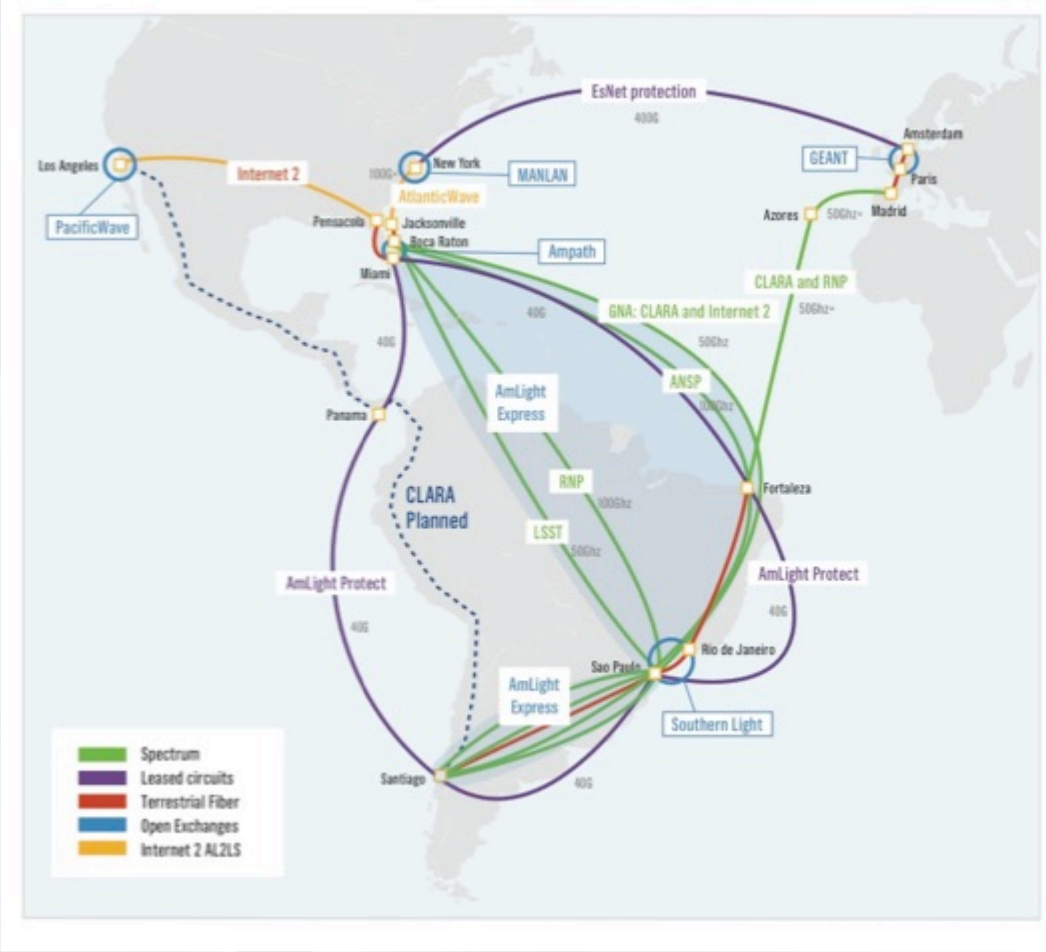

Network Virtualization and SDN Applications

- AmLight became an SDN network in 2014
- Researchers can use slicing to prototype network-aware applications
- Network Slices:
  - Defined by a set of Interfaces and VLANs
  - Each Slice has its own Openflow Controller
  - Different Topologies Available
- Slicing isolates production traffic from experimental network testbeds
  - FSFW acts as a proxy between the Data Plane
    - represented by OpenFlow devices and links, and the
  - Control Plane
    - represented by SDN applications

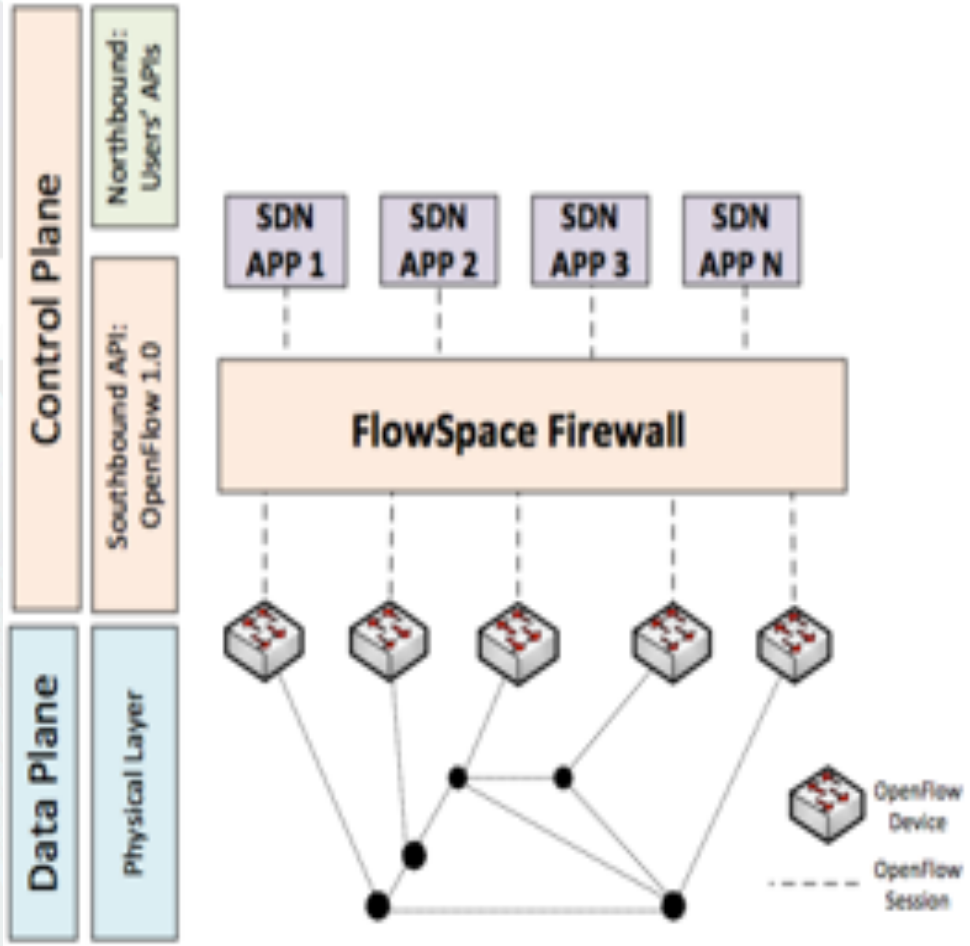

AtlanticWave-SDX Project

- AtlanticWave-SDX (Awave-SDX) is building a distributed intercontinental experimental SDX in response to a growing demand to:
  - Support end-to-end services capable of
    - Spanning multiple SDN domains
    - Dynamic provisioning of end-to-end L2 circuits
    - Providing network programmability
  - Provide more intelligent network services to
    - Foster innovation
    - Increase network efficiency
- Florida International University (FIU) and Georgia Institute of Technology (GT) are implementing AtlanticWave-SDX, in collaboration with other exchange points supporting SDN

AtlanticWave-SDX Design

- Extending SDX concept to a multi-domain SDX across 3 exchange points
- Option 1 (yellow): Single SDX controller that manages the IXP switch fabric.
- Option 2 (green): Intermediate slice manager (FlowVisor or FSFW)
  - Individual controllers assigned a slice of network resources to be managed.
  - Achieves resource isolation
- Option 3 (blue): Creates a hierarchy of controllers.
  - Local controller at each XP managed by a higher level controller
- SDX Virtual Fabric refers to a slice (a set of network headers) that will define the forwarding behavior between all the exchange points

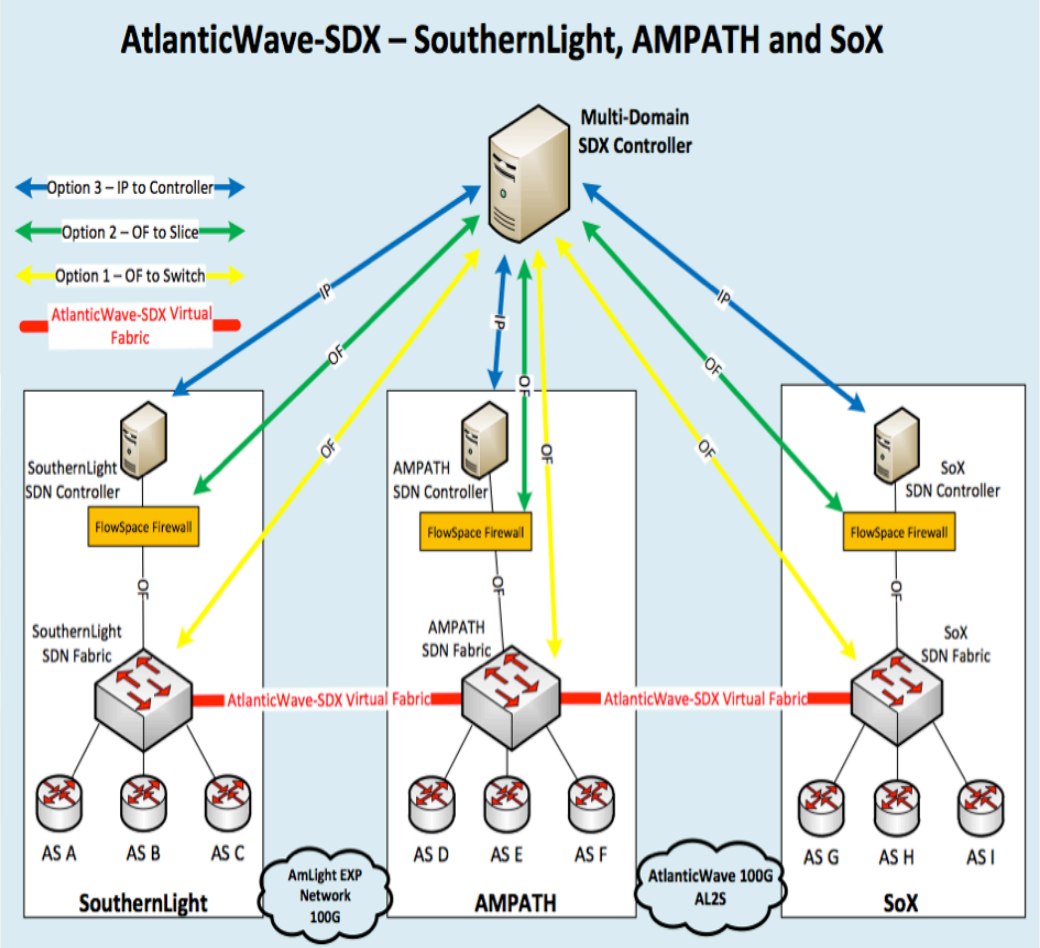

Supporting SDN applications

- Developing a software toolkit with APIs
- Application developers will be able to introduce demands to the controller
  - What resources, at what times, and with what performance requirements
- The Controller will then plan/schedule the use of resources
  - With prior knowledge of “what” and “when”
- SDX Controller Design Description and SDX specification documents have been written
- A prototype of the AtlanticWave-SDX controller is expected by Q4 2016 (I2 TechEx)

Supporting SDN Applications (2)

- Types of applications AtlanticWave-SDX will support:
  - Application-specific peering
    - Two neighbor AS exchange traffic only for certain applications
  - Inbound traffic engineering
    - AS can control inbound traffic on source IP or port number
  - Wide-area load balancing
    - Content provider can improve efficiency for client requests
  - Redirection through middle boxes
    - Redirecting targeted subsets of traffic (eg., suspected DoS traffic) through middle boxes (eg., firewall, load balancer, NAT)
  - On-demand virtual circuit provisioning
    - Provisioning virtual circuits on demand; eg., I2 AL2S, or Esnet OSCARS
  - Bandwidth Calendaring
    - Reserving bandwidth, scheduled for particular times, for an application

Science Driver: LSST

- Large-aperture wide-field ground-based 8.4 meter optical telescope
- Will take 6.4GB image every 17 seconds
- Network must transport images to NCSA within 5 seconds
- LSST operation will be composed to 2 channels:
  - A control channel and a data channel
  - Control channel handles remote operation of the telescopes by NOAO in Tucson, AZ
    - Control channel must be secure, with low latency, high priority and low bandwidth
  - Data channel will transmit 6.4GB images within 5 seconds
    - Data Channel requires high bandwidth availability, low latency and high priority
    - Support bursts close to 90G

LSST Applications and SDX

- Applications using Control Channel
  - Provision on-demand multi-domain virtual circuits with
    - Low latency, encryption, high priority and low bandwidth requirements
- Applications using Data Channel
  - Provision on-demand multi-domain virtual circuits with
    - High bandwidth availability, low latency and high priority
  - SDX reserves bandwidth for LSST to schedule image transfers

"AtlanticWave-SDX: An International SDX to Support Science Data Applications." Joaquín Chung, Jacob Cox, Julio Ibarra, Jerônimo Bezerra, Heidi Morgan, Russell Clark, Henry Owen. Supercomputing'15. November 2015

Thank You!

- NSF OpenWave, AmLight, OSDC-PIRE, CC-NIE, CC*IIE, AMPATH, AtlanticWave infrastructure, science application support, education, outreach and community building efforts are made possible by funding and support from:
  - National Science Foundation (NSF) awards ACI-1451018, ACI-1451024, ACI-1440728, ACI-0963053, ACI-1140833, ACI-1246185, ACI-1341895
  - FAPESP, ANSP – grant no. 2008/52885-8
  - Rede Nacional de Ensino e Pesquisa (RNP)
  - Association of Universities for Research in Astronomy (AURA)
  - Florida International University
  - Latin American Research and Education community
  - The many national and international collaborators who support our efforts

AmLight ExP
Americas Lightpaths Express & Protect

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

More Slides Follow

Experimental Testbeds on AmLight-ExP

- 6 network testbeds hosted at AmLight in parallel with production applications
- FIBRE Testbed
  - In partnership with RNP, a FIBRE (Future Internet testbeds / experimentation between BRazil and Europe) island was installed at AMPATH
  - AmLight was used to connect to other islands in Brazil.
  - Testbed built to
    - *Understand* what kind of requirements overlay SDN applications have and
    - *What* should be changed on production transport networks to support them

Experimental Testbeds on AmLight-ExP (2)

- ONOS SDN-IP Testbed
  - Handles BGP feeds, IP and IPv6 forwarding
  - Supports controller failover, improving resilience to SDN environments
  - A global SDN testbed without using legacy IP routers was created using OpenFlow 1.0
    - IPv6 will be supportted with OpenFlow 1.3
  - AmLight is now hosting ONOS in a production network environment
    - FIU, ANSP, RedClara, REUNA, RNP, CKLN, AARNET, GEANT/GTS, KREONET/Korea and NCTU/Taiwan are connected to the AmLight ONOS testbed

AmLight's ONOS Testbed

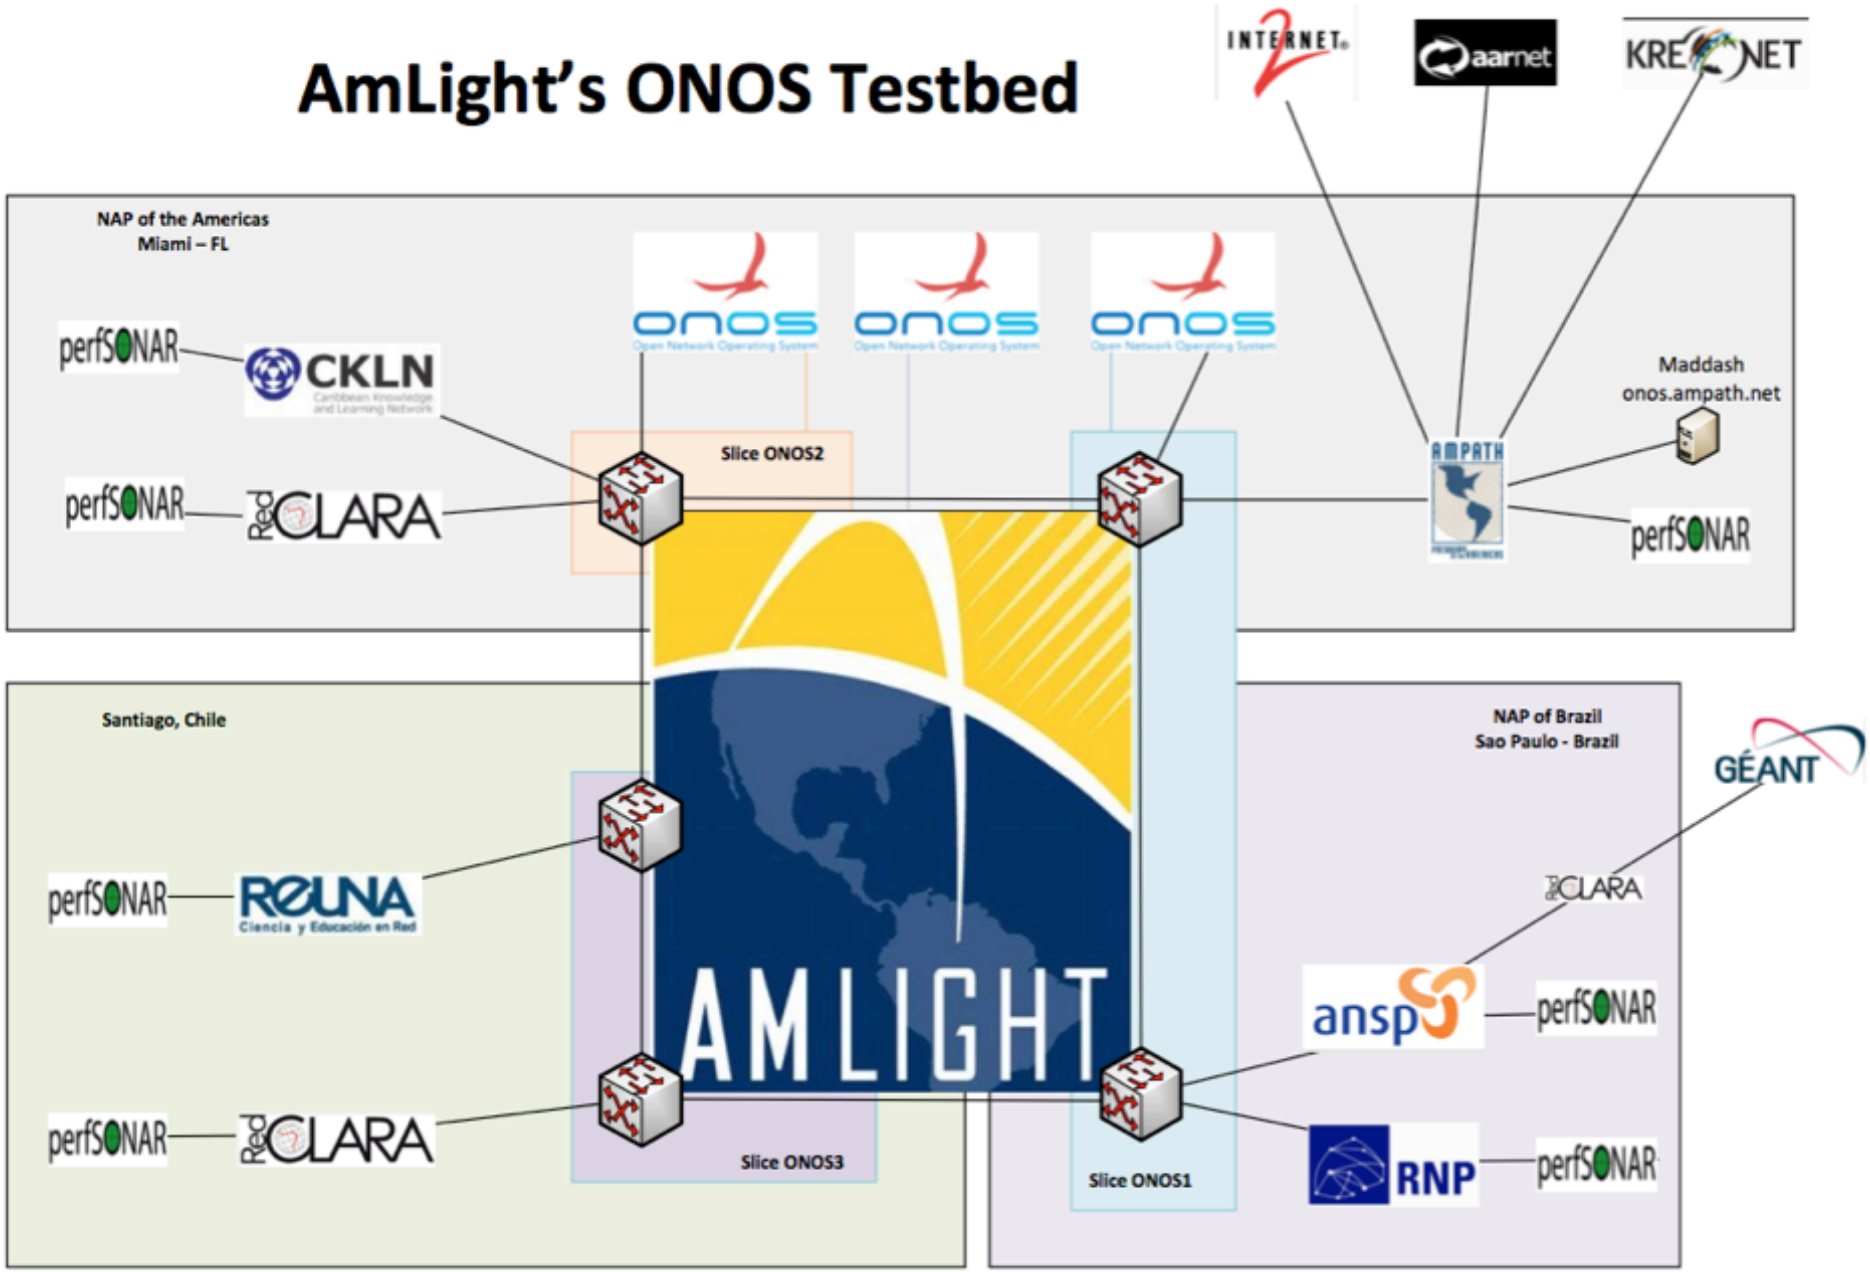

AmLight ExP
Americas Lightpaths Express & Protect

Process to Add a Slice at AmLight-ExP

- *Testbeds and production traffic share the same network infrastructure!*

- A process was created to handle slice requests:

1. User requests a slice through any channel (e-mail, phone, talk, etc.);
2. User provides info about his expectations and requirements;
3. AmLight provides a slice and servers in the *testing environment* for evaluation:
   a. User's Application is tested with the same production devices and FSFW
   b. Joint work between AmLight Engineers and User
4. Once the SDN application is "safe", both AmLight and User will manage the application in the production network:
   a. Risky but less manpower required

Programmability @ AmLight-ExP

Two possible interfaces to use AmLight SDN offered to users and researchers:

OpenFlow (currently 1.0)

- Dedicated slices/Dedicated VLAN range
- Different virtual topologies available
- Layer 2 and 3 matches
- Low level configuration

NSI v2 – Network Service Interface

- High level abstraction for layer 2 multi-domain provisioning
- No need to know the topology and physical devices/configurations
- Layer 2 circuit provided as a service: easier to isolate from production traffic

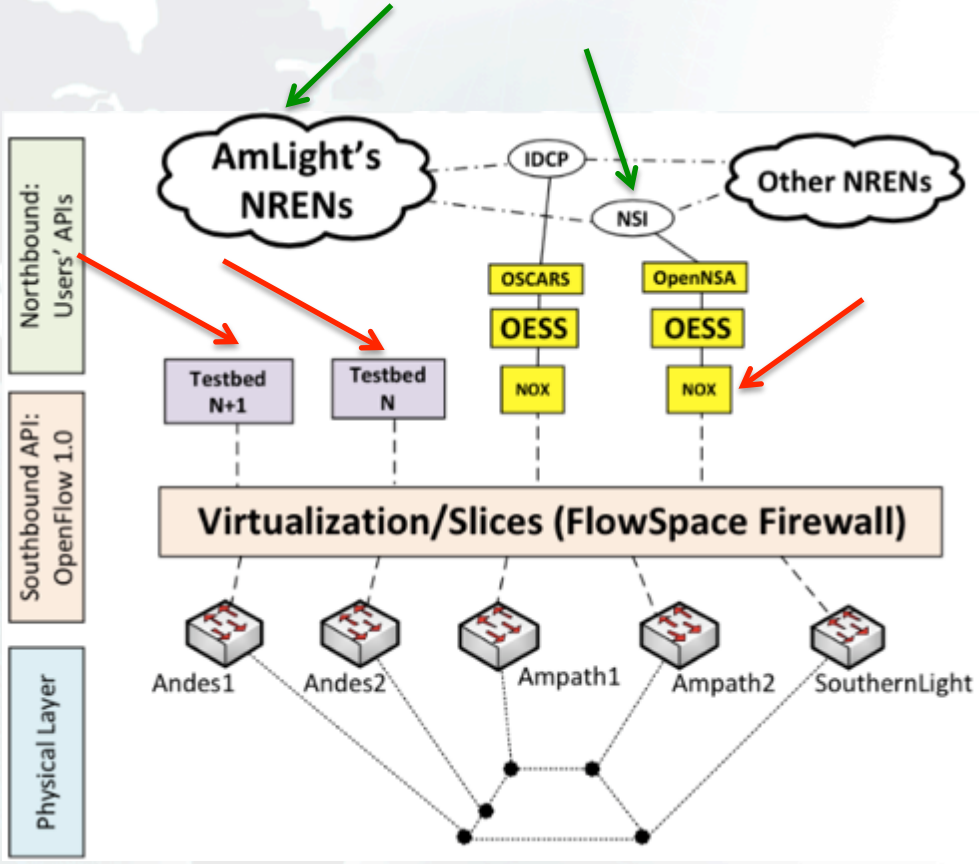

Future: Developing a higher-level interface as powerful as OpenFlow

Future SDN support at AmLight-ExP

- AtlanticWave-SDX:
  - Multi domain SDX controller
  - High Level API
  - Support end-to-end services capable of
    - Dynamic provisioning of end-to-end L2/L3/L4 circuits
  - Application developers will be able to introduce demands to the controller
    - What resources, at what times, and with what performance requirements
  - The Controller will then plan/schedule the use of resources
    - With prior knowledge of “what” and “when”
- A prototype of the AtlanticWave-SDX controller will be presented at the Internet2 Tech Exchange

Network Infrastructure Design

- Years 1 and 2:
  - Upgrade AMPATH IXP infrastructure to support 140G in year 1
  - Deploy new technologies at AMPATH to fully support SDN in its switching fabric
- Years 3, 4 and 5:
  - Upgrade the switching capacity at AMPATH to receive 6 100G links from AmLight ExP
  - Extend capacity to Jacksonville over the FLR network to the Internet2 AL2S
  - FLR providing two sets of 250GHz channels in its backbone, provisioned over diverse paths